

US009684570B1

(12) **United States Patent**
Wilding

(10) **Patent No.:** **US 9,684,570 B1**
(45) **Date of Patent:** **Jun. 20, 2017**

(54) **SYSTEMS, METHODS, AND APPARATUSES FOR FIXING LOGICAL OR PHYSICAL CORRUPTION IN DATABASES USING IMMUTABLE LSM TREES**

(71) Applicant: **Mark Wilding**, Issaquah, WA (US)

(72) Inventor: **Mark Wilding**, Issaquah, WA (US)

(73) Assignee: **salesforce.com, inc.**, San Francisco, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: **14/135,564**

(22) Filed: **Dec. 19, 2013**

Related U.S. Application Data

(60) Provisional application No. 61/739,109, filed on Dec. 19, 2012.

(51) **Int. Cl.**
G06F 11/00 (2006.01)
G06F 11/14 (2006.01)
H04L 29/08 (2006.01)
G06F 17/30 (2006.01)

(52) **U.S. Cl.**
CPC **G06F 11/1471** (2013.01); **G06F 11/1474** (2013.01); **G06F 17/30371** (2013.01); **H04L 67/10** (2013.01); **G06F 2201/80** (2013.01); **G06F 2201/805** (2013.01)

(58) **Field of Classification Search**
CPC **G06F 11/1471**; **G06F 11/1474**; **G06F 17/30286**; **G06F 17/30371**
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

8,977,898 B1* 3/2015 Veeraswamy G06F 17/30227
714/15
2012/0011106 A1* 1/2012 Reid G06F 9/466
707/695

OTHER PUBLICATIONS

O'Neil et al., "The Log-Structured Merge-Tree (LSM-Tree)", 1996, pp. 1-32.*

* cited by examiner

Primary Examiner — Michael Maskulinski

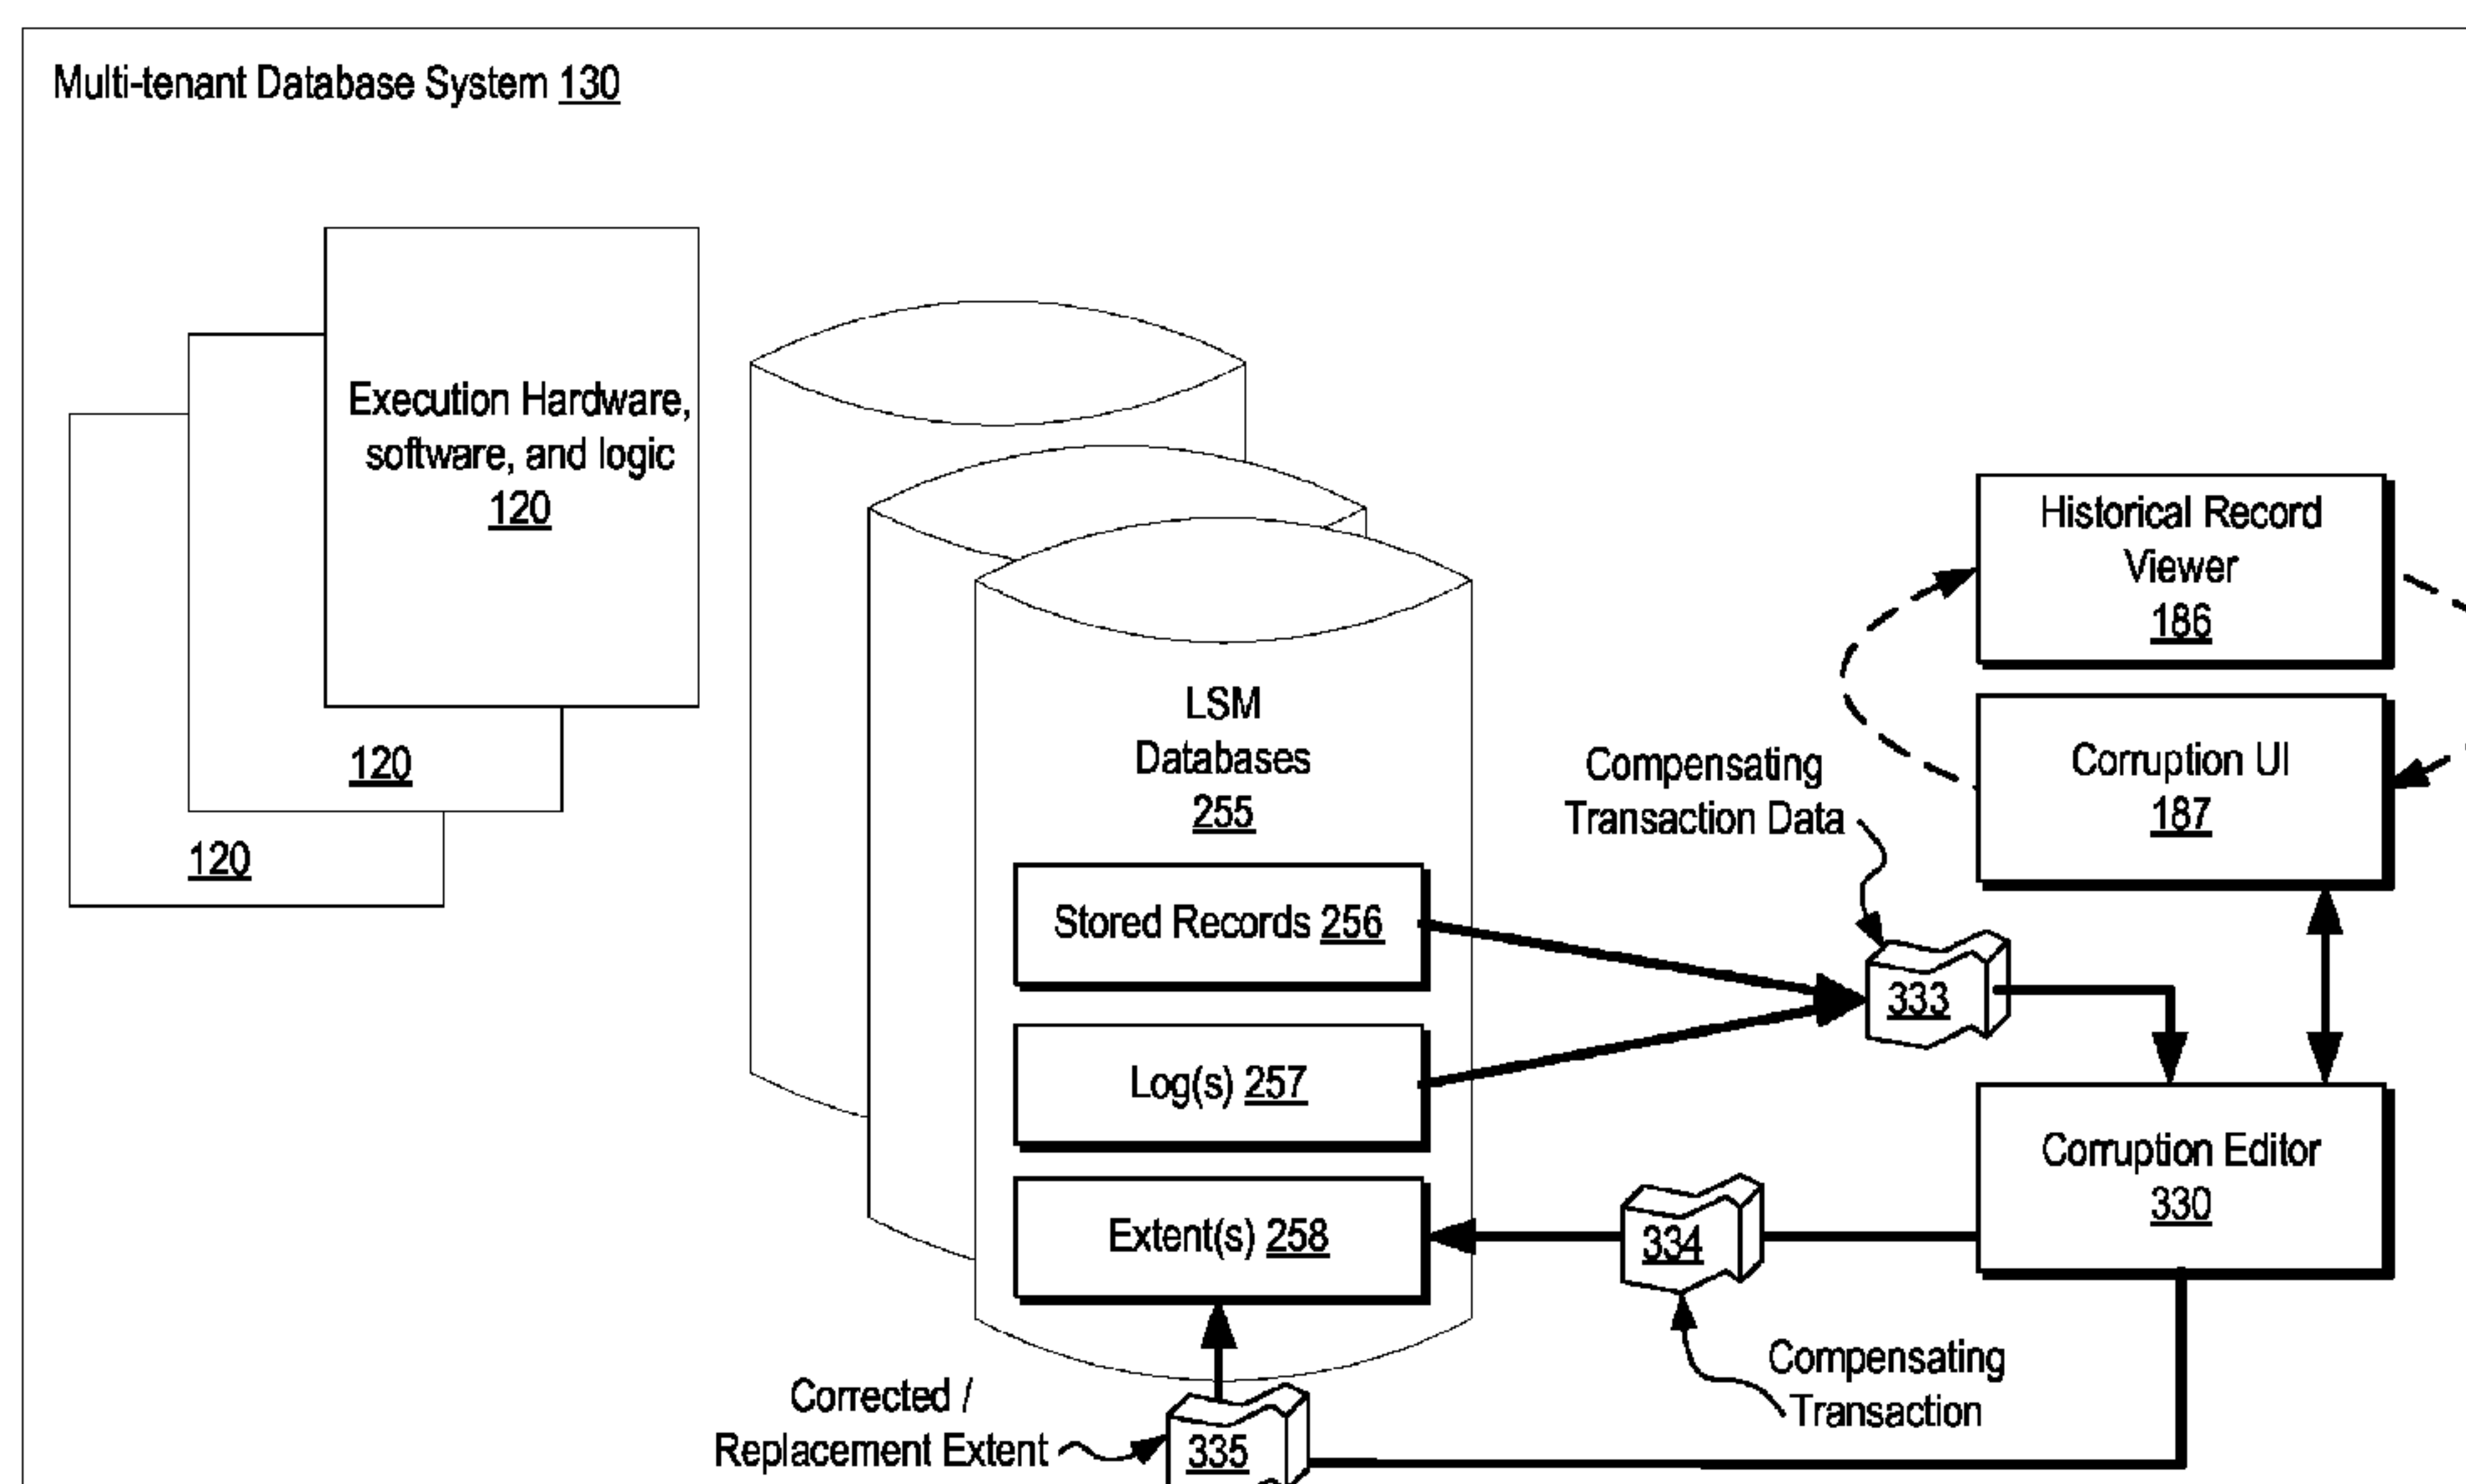
(74) *Attorney, Agent, or Firm* — Blakely Sokoloff Taylor & Zafman LLP

(57) **ABSTRACT**

In accordance with disclosed embodiments, there are provided methods, systems, and apparatuses for fixing logical or physical corruption in databases using immutable LSM trees including, for example, means for storing records in a LSM database at the host organization; processing transactions to the LSM database by writing the transactions into immutable extents for subsequent merge into the LSM database; recording a history of the transactions in a log; identifying corruption in at least one of the records of the LSM database represented within the immutable extents prior to merge of the immutable extents to the LSM database; and correcting the identified corruption to the at least one of the records of the LSM database based at least in part on the recorded history of the transactions in the log during live operation of the LSM database. Other related embodiments are disclosed.

23 Claims, 11 Drawing Sheets

↪ 300






FIG. 2
200





FIG. 3
300



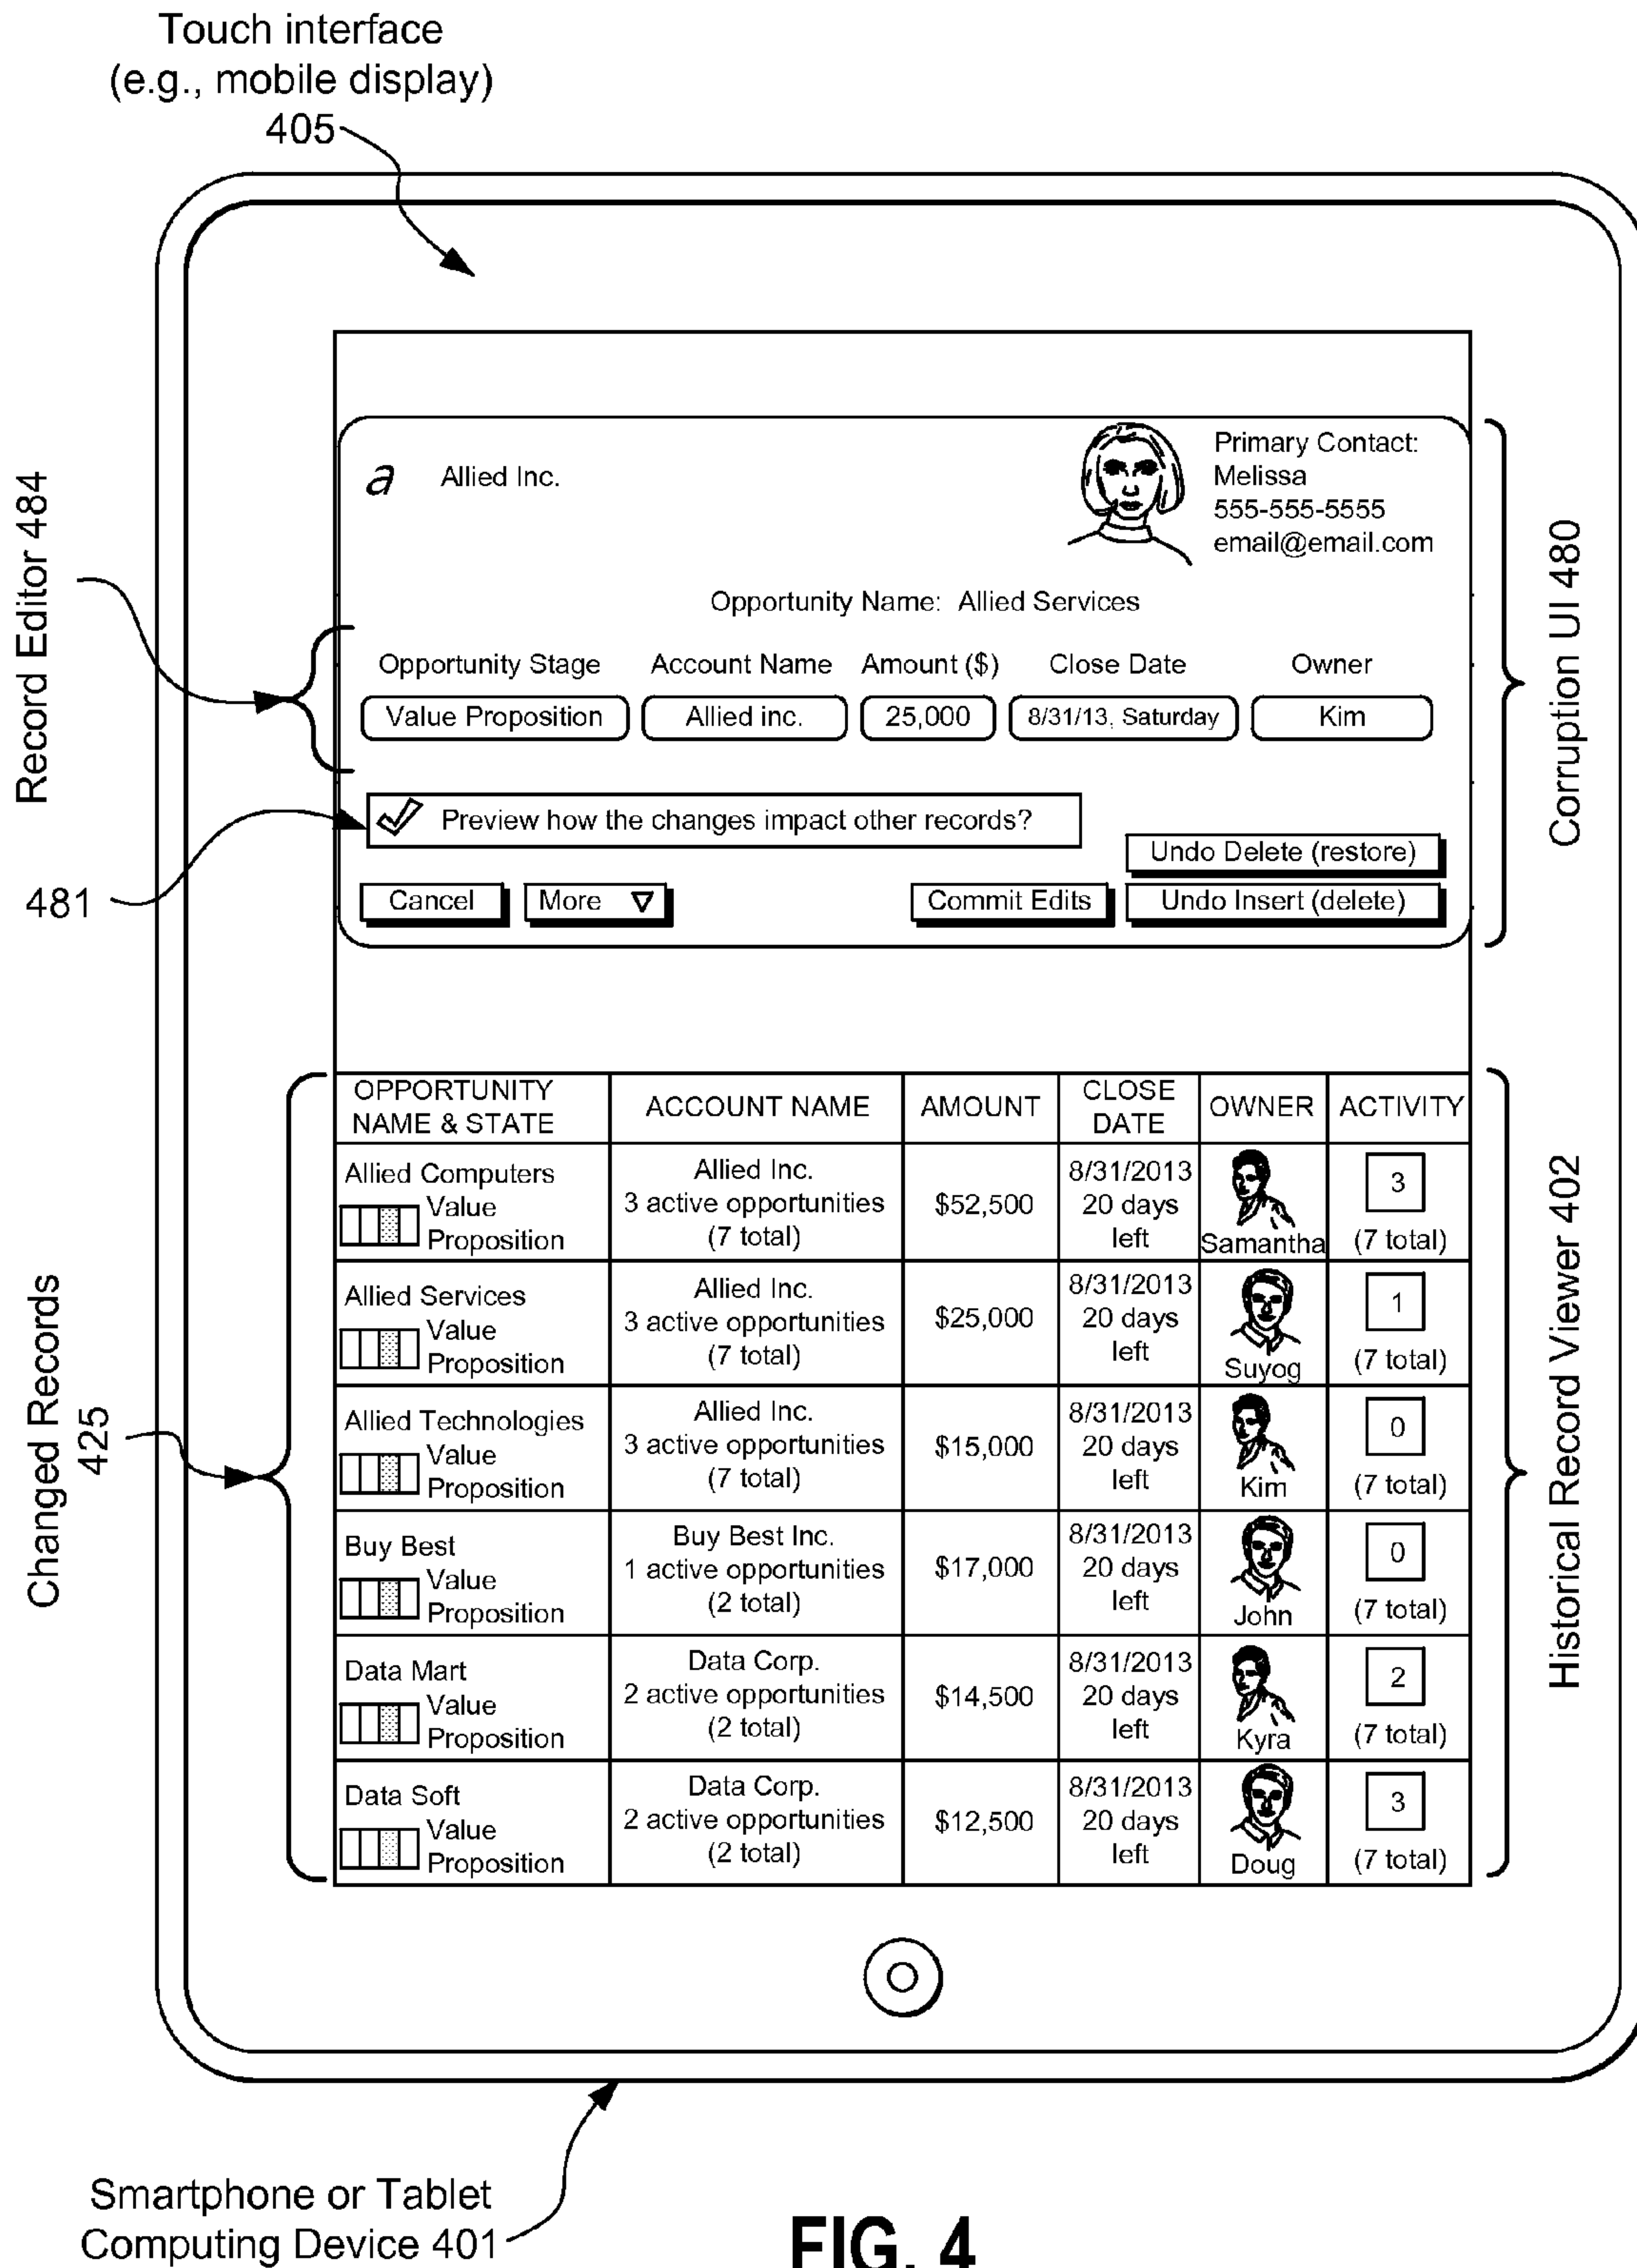


FIG. 4

FIG. 5A




FIG. 5B
500

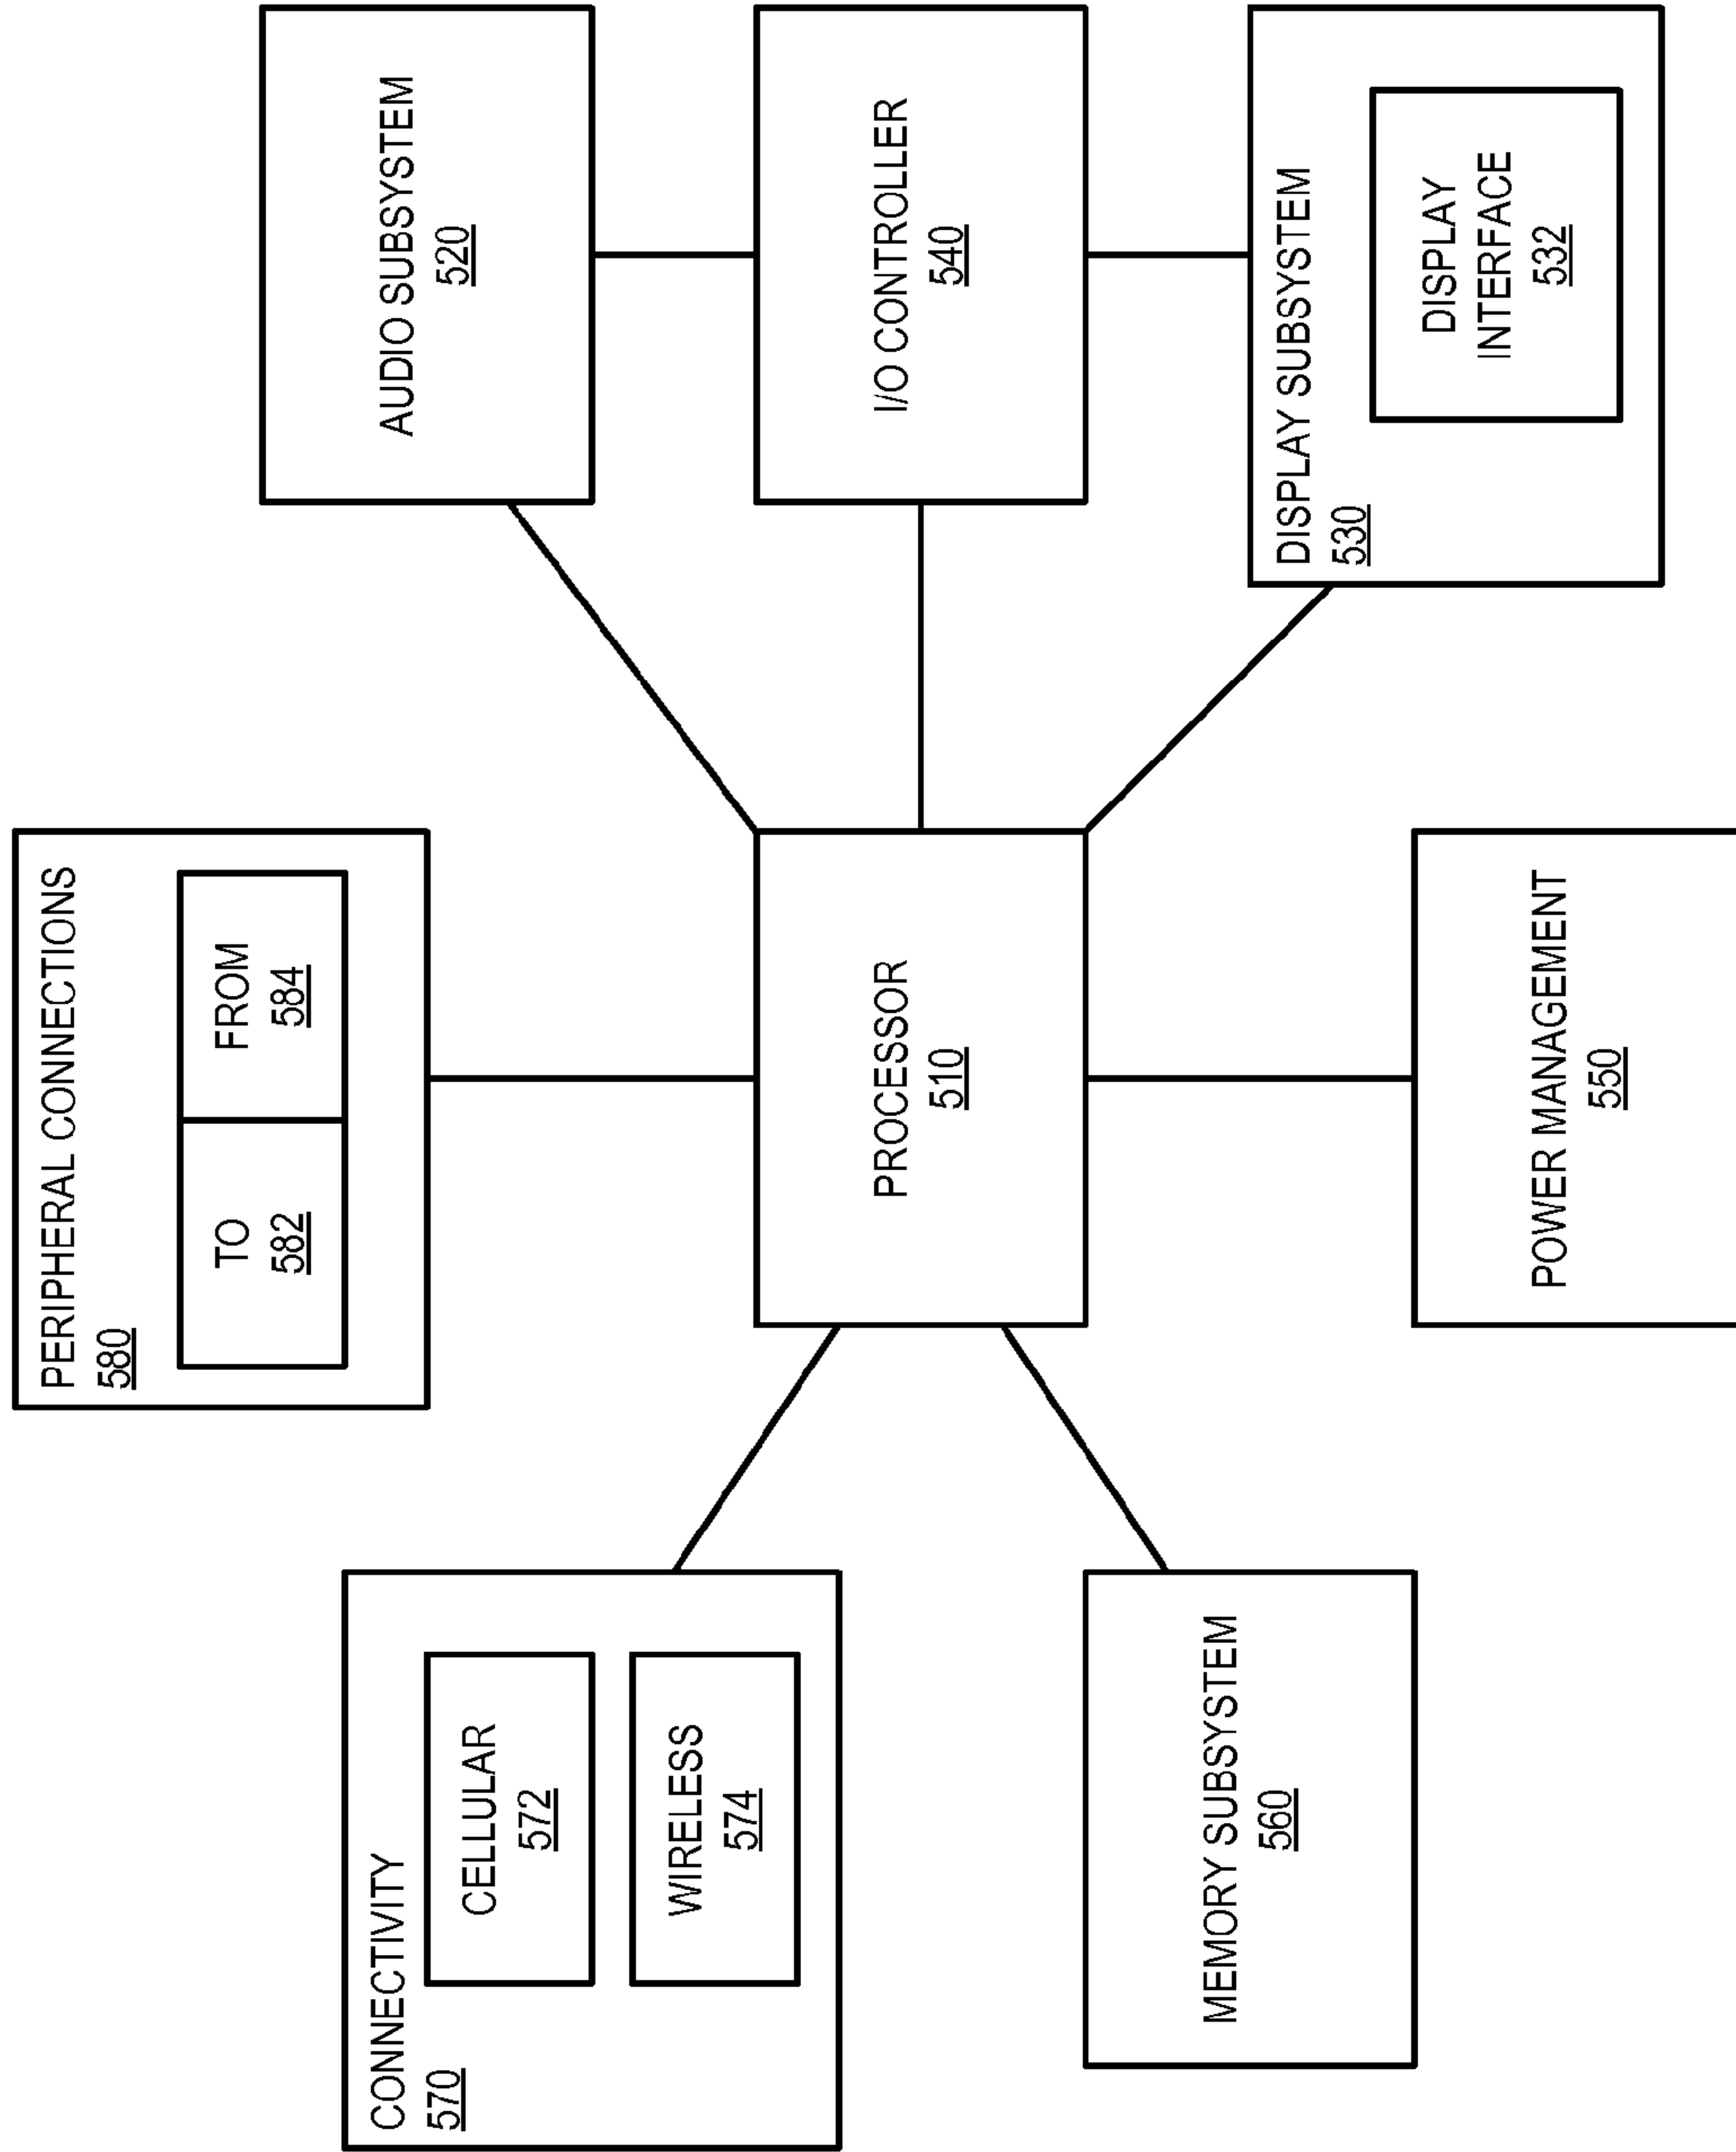
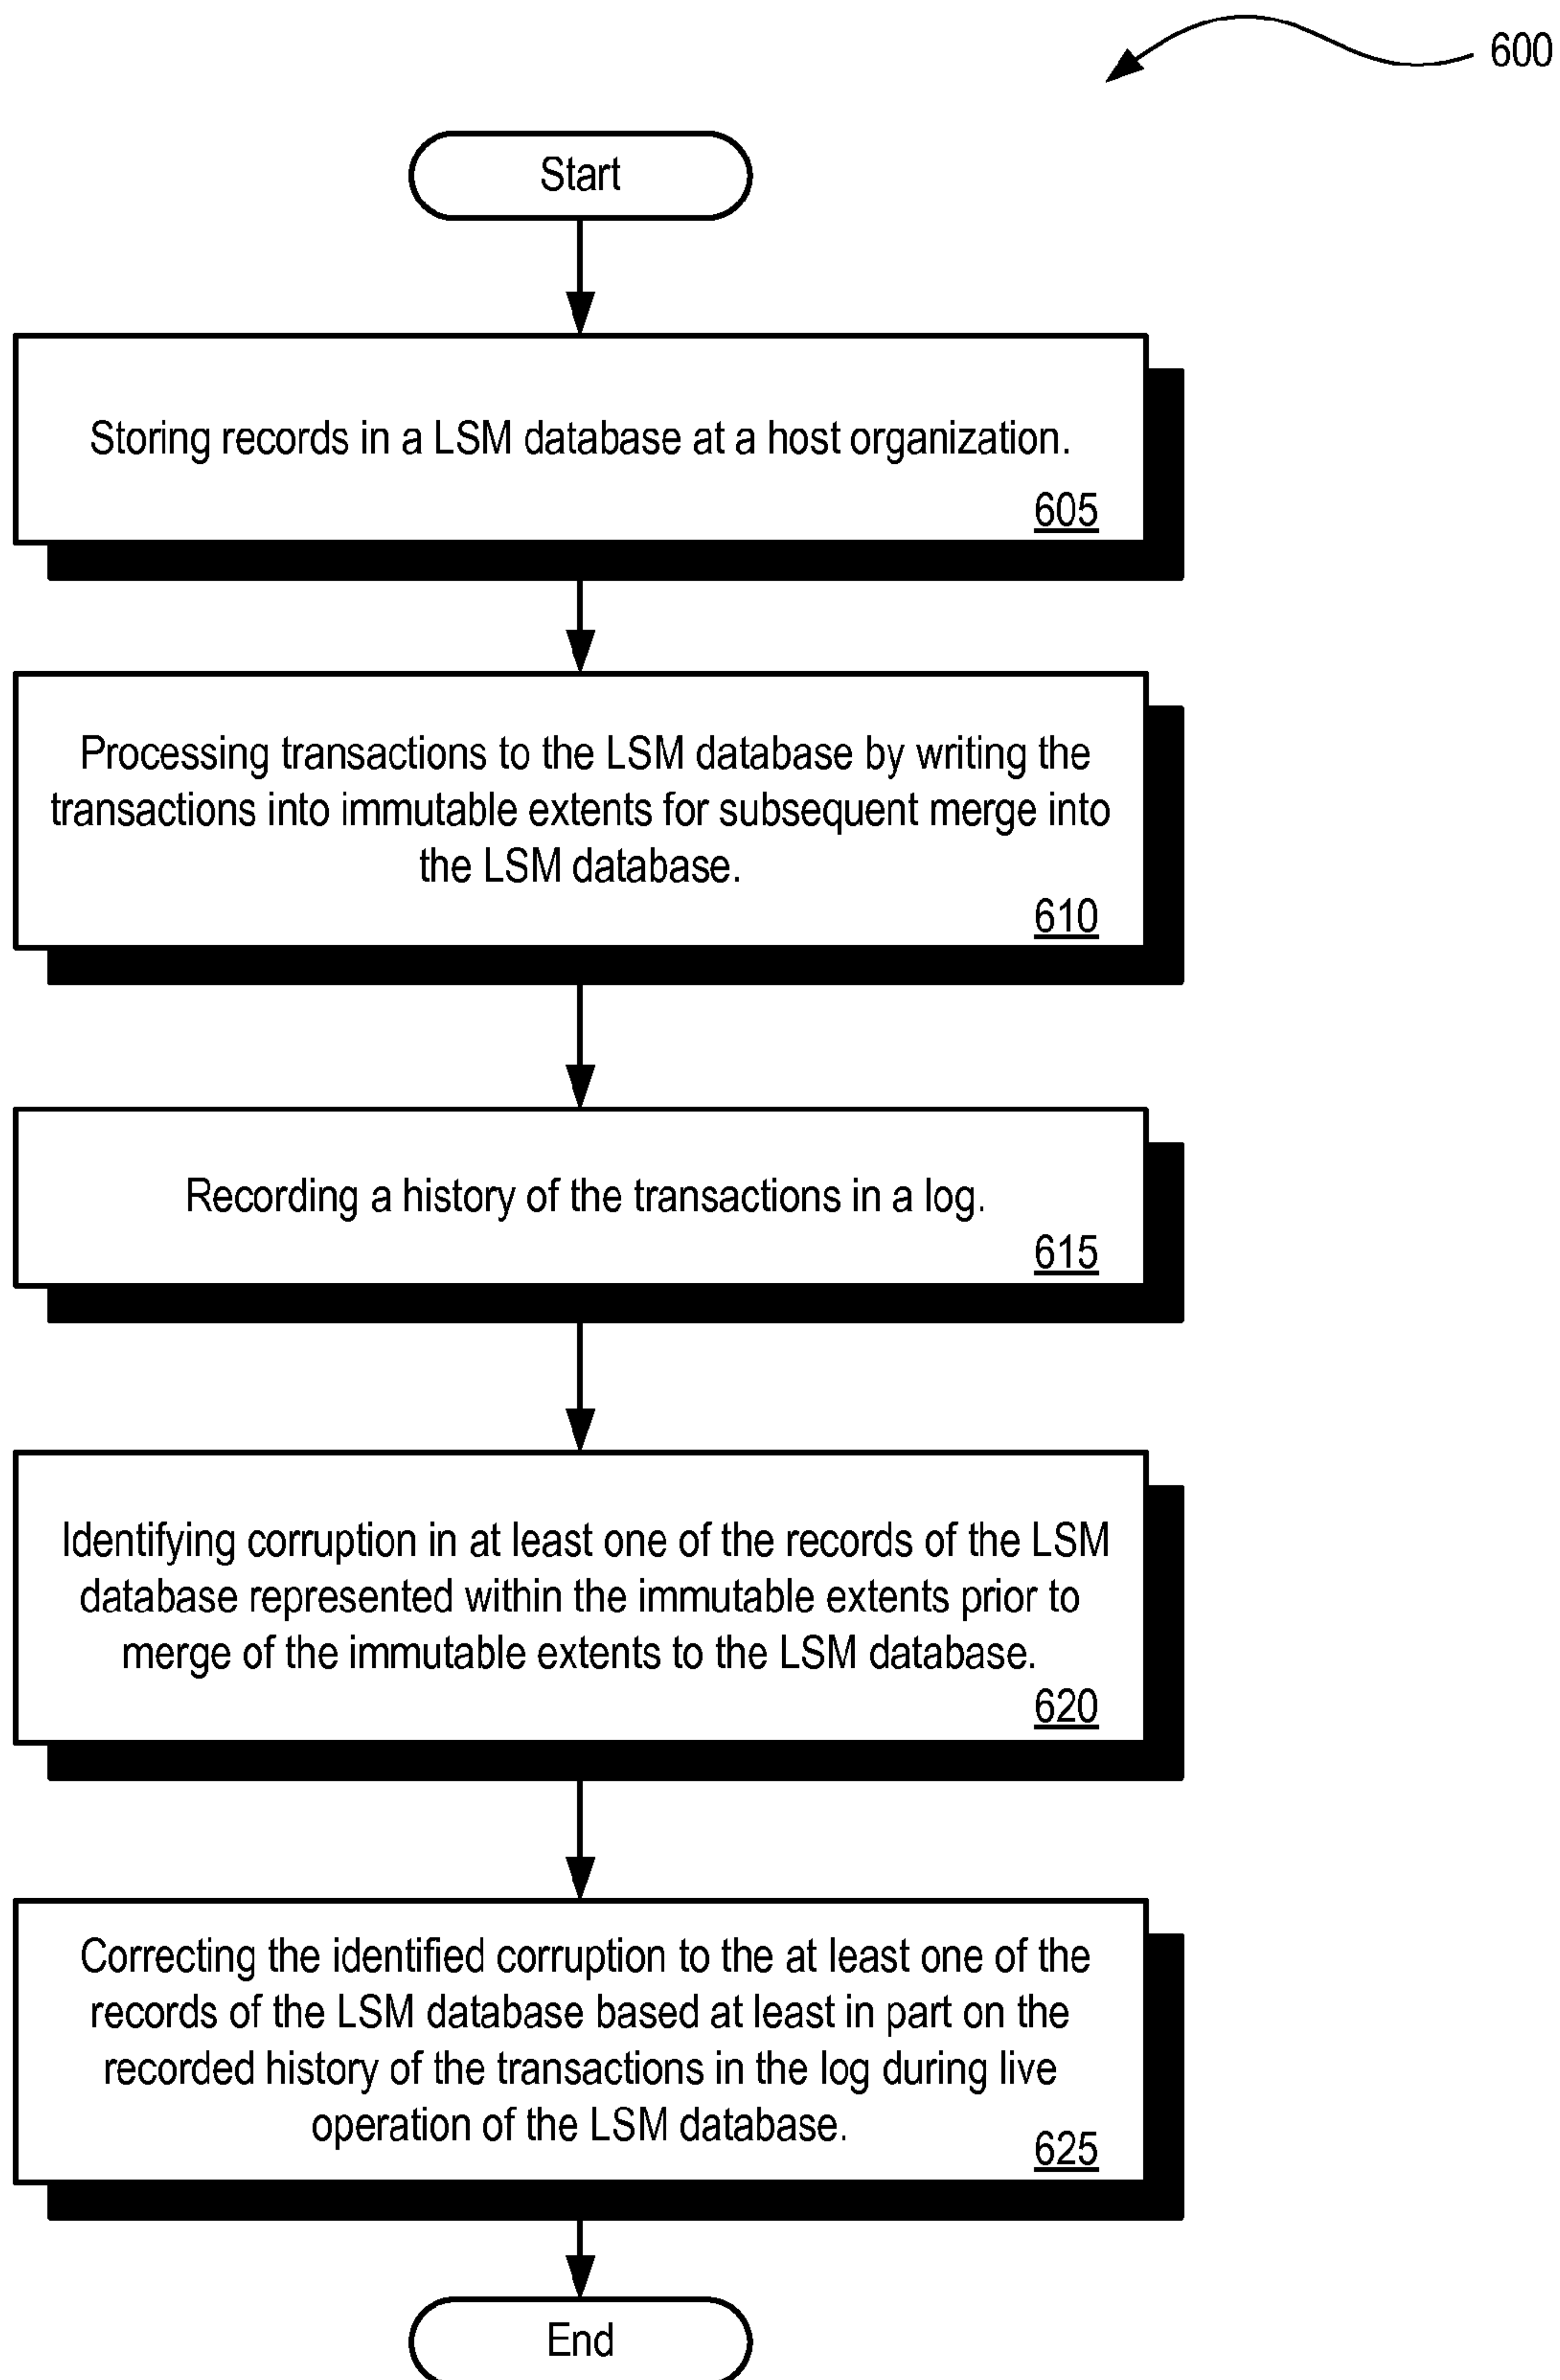


FIG. 6



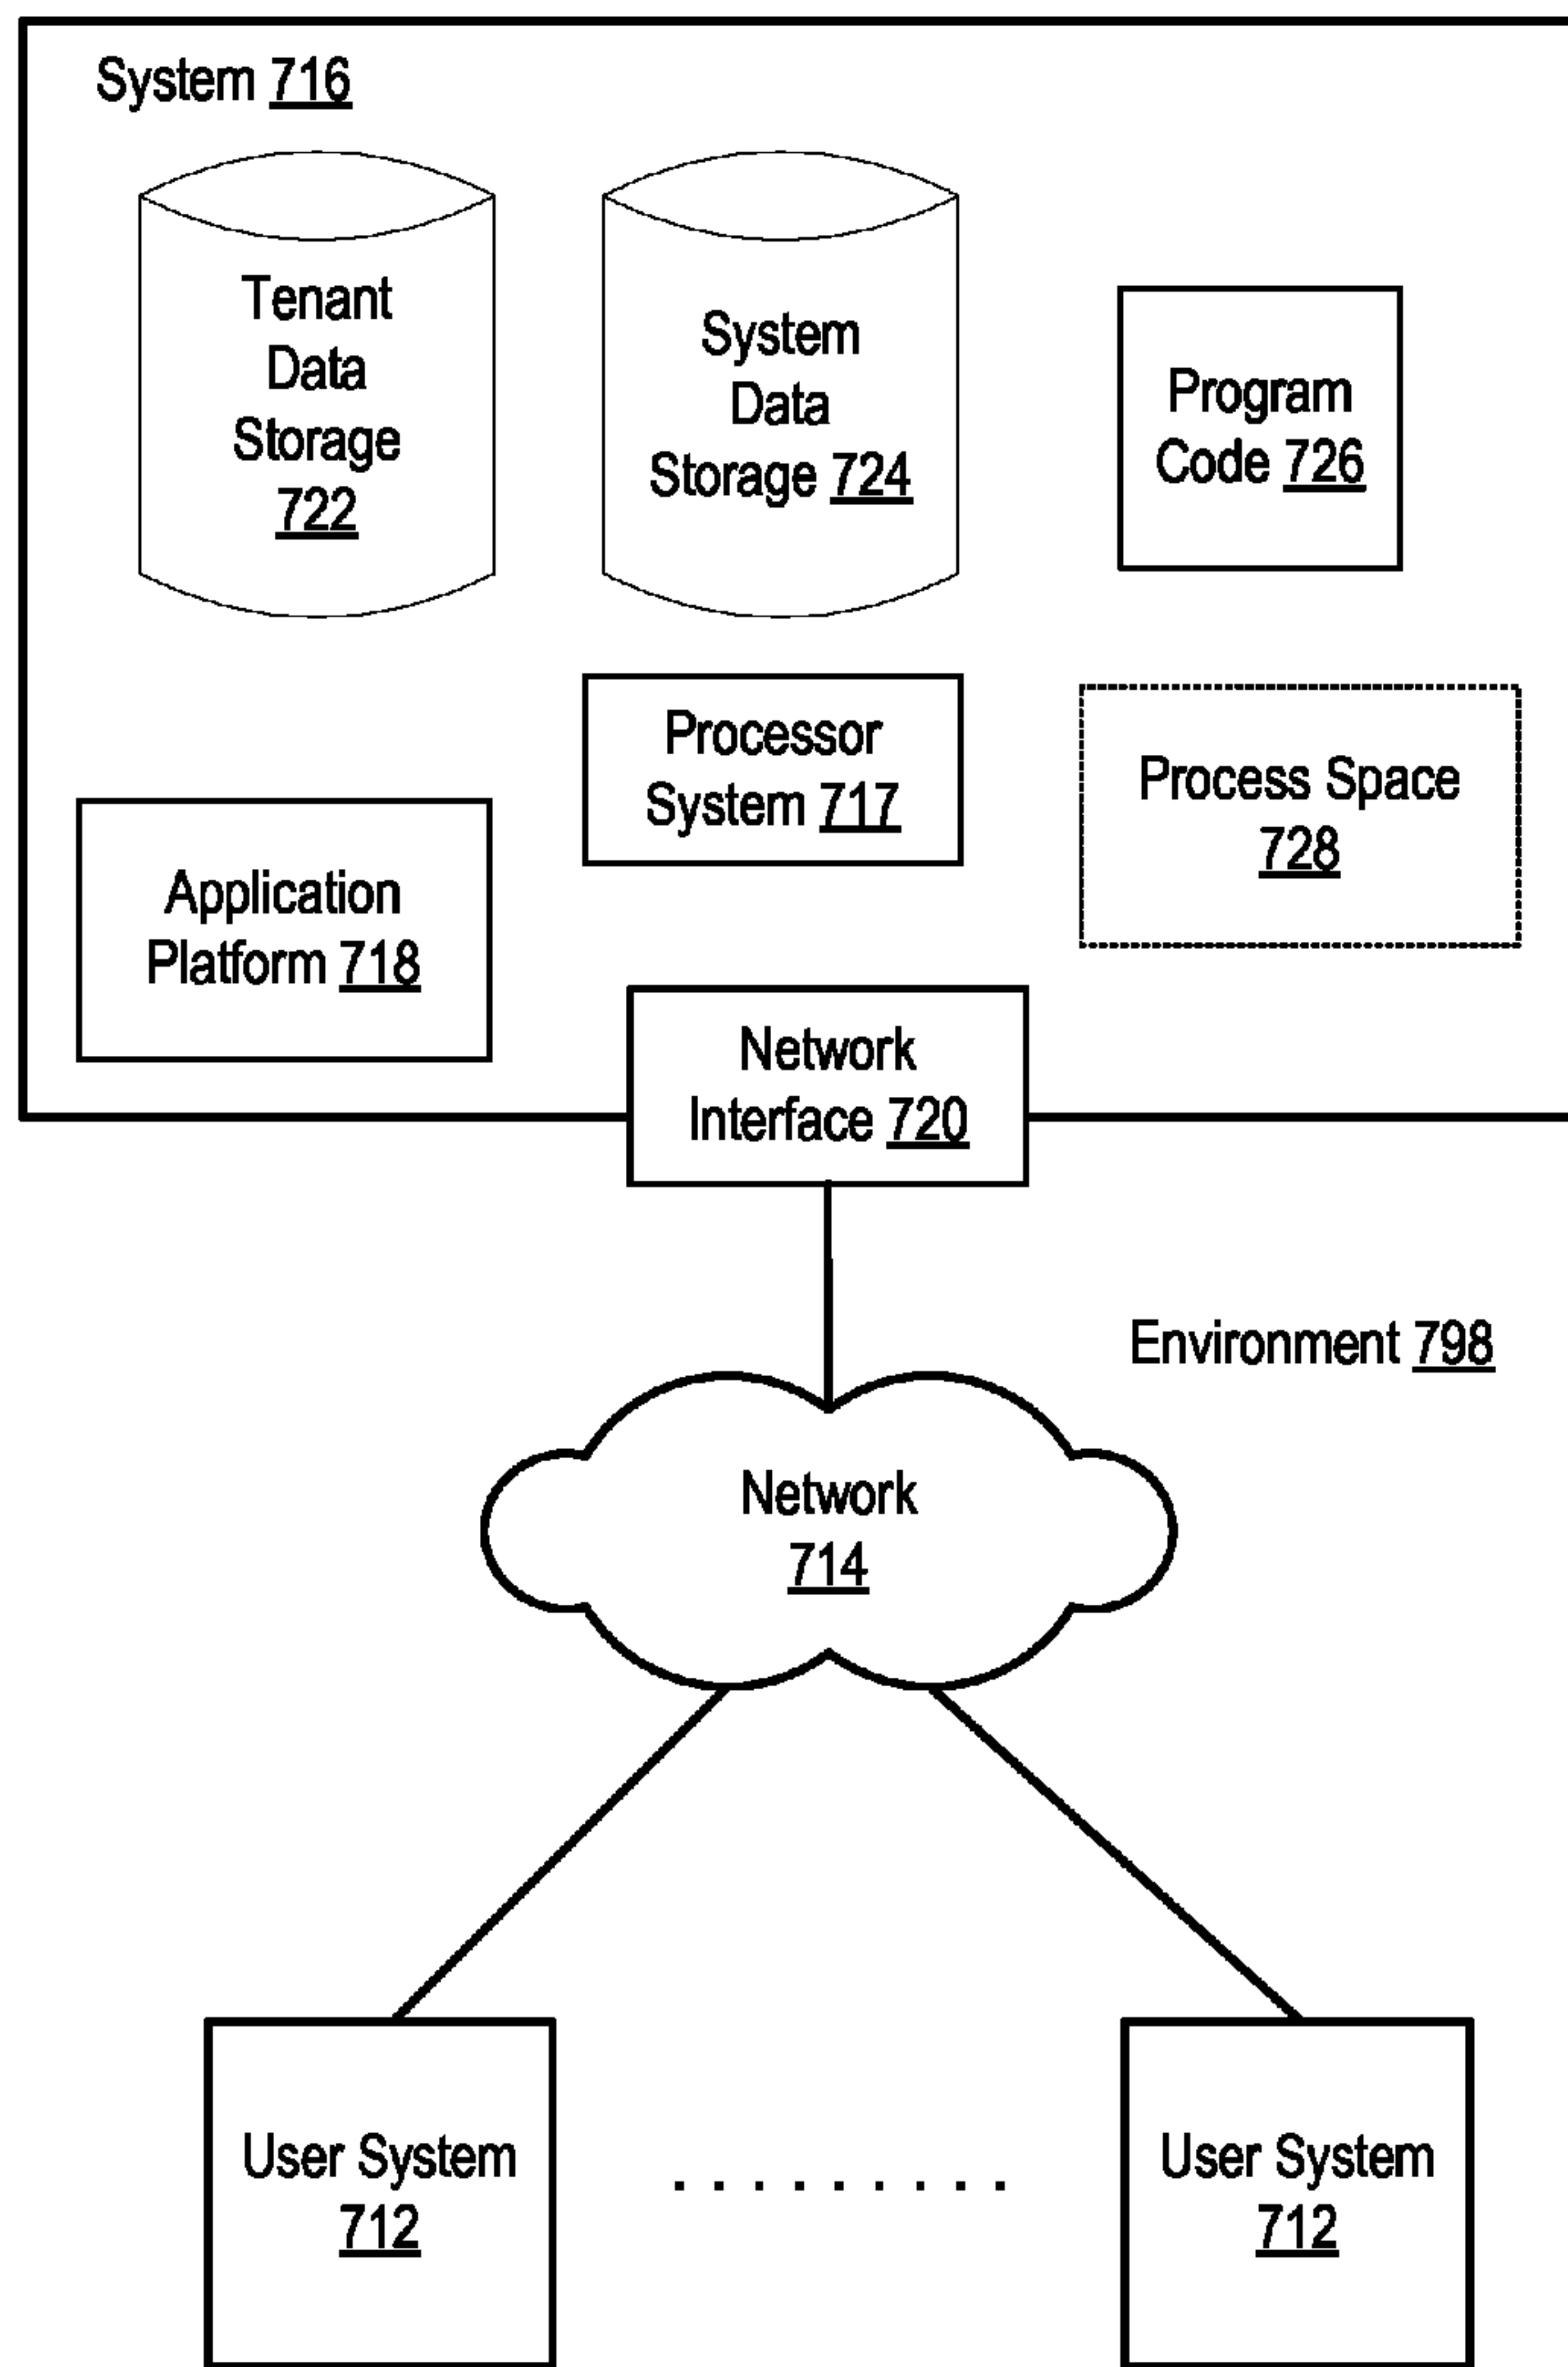


FIG. 7A





FIG. 7B






FIG. 9

1

**SYSTEMS, METHODS, AND APPARATUSES
FOR FIXING LOGICAL OR PHYSICAL
CORRUPTION IN DATABASES USING
IMMUTABLE LSM TREES**

CLAIM OF PRIORITY

This application is related to, and claims priority to, the provisional utility application entitled "FIXING LOGICAL OR PHYSICAL CORRUPTION WITH IMMUTABLE LSM TREES," filed on Dec. 19, 2012, having an application No. of 61/739,109, the entire contents of which are incorporated herein by reference.

COPYRIGHT NOTICE

A portion of the disclosure of this patent document contains material which is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure, as it appears in the Patent and Trademark Office patent file or records, but otherwise reserves all copyright rights whatsoever.

TECHNICAL FIELD

Embodiments of the invention relate generally to the field of computing, and more particularly, to systems, methods, and apparatuses for fixing logical or physical corruption in databases using immutable LSM trees.

BACKGROUND

The subject matter discussed in the background section should not be assumed to be prior art merely as a result of its mention in the background section. Similarly, a problem mentioned in the background section or associated with the subject matter of the background section should not be assumed to have been previously recognized in the prior art. The subject matter in the background section merely represents different approaches, which in and of themselves may also correspond to embodiments of the claimed inventions.

A single a multi-tenant database system operates to store data on behalf of a multitude of paying subscribers, each being a "tenant" of the database system, hence the term multi-tenant database system.

Within such an operational environment, computational efficiency, system responsiveness, and data integrity are all of paramount concern both to the provider of the multi-tenant database system and to the subscribers or tenants of such a system. Moreover, with on-demand technologies having multiple distinct clients simultaneously utilizing the system and relying upon its availability it is critical to avoid service outages which can create frustration on behalf of users, create bad press coverage for the host organization, and damage the credibility of the host organization which in turn creates a negative business environment making it more difficult to acquire new customers and retain existing customers.

Within conventional database systems, data corruption detected by the database engine software is considered a catastrophic event, as it should be, causing the database to very often "crash" rather than risk serving corrupted data in reply to queries. Some enterprise level databases do not crash and perform special isolation techniques instead. Nevertheless, a crashed database may be acceptable in a single tenant environment where one entity hosts their data on the

2

database and is responsible for maintaining their own database internal to an organization because a catastrophic failure and database crash will result in only that particular tenant's users being affected.

5 However, consider the environment in which a multi-tenant database system operates as an on-demand or cloud based subscription service providing database services to tens of thousands of customers. A catastrophic failure in such an environment due to data corruption will affect a large number of customer organizations, their business operations, their users, customers of those businesses, and so forth. Even if a catastrophic failure is limited within a multi-tenant database system to a single server pod or some logical subset, the failure will still affect hundreds or thousands of customers having data on such a sub-set rather than only a single entity as would occur in a single-tenant database system. In addition to the bad press and loss of credibility, such a database crash can directly lead to millions of dollars per hour of lost revenue for the duration of the outage. Further still, a single tenant of a multi-tenant database system cannot be permitted to trigger a database outage that could affect potentially thousands of other tenants of the same multi-tenant database system, each of whom are running their own businesses. It would be grossly unfair for a single tenant to detrimentally impact so many others in such a way.

Rectifying database corruption is additionally notoriously difficult and requires skilled database experts and technicians, which in turn leads to further cost and delay in the recovery of database services. At the same time, corruption in a database is a critical problem and cannot simply be ignored as returning corrupted data in reply to queries could have even more damaging results than returning no data at all due to a service outage.

The present state of the art may therefore benefit from the systems, methods, and apparatuses for fixing logical or physical corruption in databases using immutable LSM trees as described herein.

BRIEF DESCRIPTION OF THE DRAWINGS

Embodiments are illustrated by way of example, and not by way of limitation, and will be more fully understood with reference to the following detailed description when considered in connection with the figures in which:

FIG. 1 depicts an exemplary architecture in accordance with described embodiments;

FIG. 2 depicts another exemplary architecture in accordance with described embodiments;

FIG. 3 depicts another exemplary architecture in accordance with described embodiments;

FIG. 4 depicts an exemplary graphical interface operating at a mobile, smartphone, or tablet computing device in accordance with the embodiments;

FIG. 5A depicts a tablet computing device and a handheld smartphone each having a circuitry integrated therein as described in accordance with the embodiments;

FIG. 5B is a block diagram of an embodiment of tablet computing device, a smart phone, or other mobile device in which touchscreen interface connectors are used; and

FIG. 6 is a flow diagram illustrating a method for fixing logical or physical corruption in databases using immutable LSM trees in accordance with disclosed embodiments;

FIG. 7A illustrates a block diagram of an environment in which an on-demand database service may operate in accordance with the described embodiments;

FIG. 7B illustrates another block diagram of an embodiment of elements of FIG. 7A and various possible interconnections between such elements in accordance with the described embodiments;

FIG. 8 illustrates a diagrammatic representation of a machine in the exemplary form of a computer system, in accordance with one embodiment; and

FIG. 9 shows a diagrammatic representation of a system in accordance with which embodiments may operate, be installed, integrated, or configured.

DETAILED DESCRIPTION

Described herein are systems, methods, and apparatuses for fixing logical or physical corruption in databases using immutable LSM trees in an on-demand service environment. Such systems, devices, and methods may include, for example, means for: storing records in a LSM database at the host organization; processing transactions to the LSM database by writing the transactions into immutable extents for subsequent merge into the LSM database; recording a history of the transactions in a log; identifying corruption in at least one of the records of the LSM database represented within the immutable extents prior to merge of the immutable extents to the LSM database; and correcting the identified corruption to the at least one of the records of the LSM database based at least in part on the recorded history of the transactions in the log during live operation of the LSM database.

The described methodologies for fixing logical or physical corruption in databases using immutable LSM trees are work and is specifically designed for corruptions in extents that are found at any time. The extents may periodically go through a merge process to create new extents. According to the described embodiments, sufficient redundancy in the LSM database is maintained so as to re-merge and therefore create any data extent at any time necessary. According to certain embodiments, sufficient log data is additionally maintained for both tenant corruption above the API/SQL interface as well as to recreate data extents higher in the LSM tree during live operation of the database. That is to say, the corruption is fixed without taking down or crashing the database.

In the following description, numerous specific details are set forth such as examples of specific systems, languages, components, etc., in order to provide a thorough understanding of the various embodiments. It will be apparent, however, to one skilled in the art that these specific details need not be employed to practice the embodiments disclosed herein. In other instances, well known materials or methods have not been described in detail in order to avoid unnecessarily obscuring the disclosed embodiments.

In addition to various hardware components depicted in the figures and described herein, embodiments further include various operations which are described below. The operations described in accordance with such embodiments may be performed by hardware components or may be embodied in machine-executable instructions, which may be used to cause a general-purpose or special-purpose processor programmed with the instructions to perform the operations. Alternatively, the operations may be performed by a combination of hardware and software.

Embodiments also relate to an apparatus for performing the operations disclosed herein. This apparatus may be specially constructed for the required purposes, or it may be a general purpose computer selectively activated or reconfigured by a computer program stored in the computer. Such

a computer program may be stored in a computer readable storage medium, such as, but not limited to, any type of disk including floppy disks, optical disks, CD-ROMs, and magnetic-optical disks, read-only memories (ROMs), random access memories (RAMs), EPROMs, EEPROMs, magnetic or optical cards, or any type of media suitable for storing electronic instructions, each coupled to a computer system bus.

The algorithms and displays presented herein are not inherently related to any particular computer or other apparatus. Various general purpose systems may be used with programs in accordance with the teachings herein, or it may prove convenient to construct more specialized apparatus to perform the required method steps. The required structure for a variety of these systems will appear as set forth in the description below. In addition, embodiments are not described with reference to any particular programming language. It will be appreciated that a variety of programming languages may be used to implement the teachings of the embodiments as described herein.

Embodiments may be provided as a computer program product, or software, that may include a machine-readable medium having stored thereon instructions, which may be used to program a computer system (or other electronic devices) to perform a process according to the disclosed embodiments. A machine-readable medium includes any mechanism for storing or transmitting information in a form readable by a machine (e.g., a computer). For example, a machine-readable (e.g., computer-readable) medium includes a machine (e.g., a computer) readable storage medium (e.g., read only memory (“ROM”), random access memory (“RAM”), magnetic disk storage media, optical storage media, flash memory devices, etc.), a machine (e.g., computer) readable transmission medium (electrical, optical, acoustical), etc.

Any of the disclosed embodiments may be used alone or together with one another in any combination. Although various embodiments may have been partially motivated by deficiencies with conventional techniques and approaches, some of which are described or alluded to within the specification, the embodiments need not necessarily address or solve any of these deficiencies, but rather, may address only some of the deficiencies, address none of the deficiencies, or be directed toward different deficiencies and problems where are not directly discussed.

FIG. 1 depicts an exemplary architecture **100** in accordance with described embodiments. In one embodiment, a production environment **111** is communicably interfaced with a plurality of client devices **106A-C** (e.g., such as mobile devices, smart phones, tablets, PCs, etc.) through host organization **110**. In one embodiment, a multi-tenant database system **130** includes databases **155**, for example, to store tables, datasets, and underlying database records with user data on behalf of customer organizations **105A-C** or users.

Multi-tenant database system **130** includes a plurality of underlying hardware, software, and logic elements **120** that implement database functionality and a code execution environment within the host organization **110**. In accordance with one embodiment, multi-tenant database system **130** further implements databases **155** to service database queries and other data interactions with the databases **155**. The hardware, software, and logic elements **120** of the multi-tenant database system **130** are separate and distinct from a plurality of customer organizations (**105A**, **105B**, and **105C**) which utilize the services provided by the host organization **110** by communicably interfacing to the host organization

110 via network 125. In such a way, host organization 110 may implement on-demand database services or cloud computing services to subscribing customer organizations 105A-C.

Host organization 110 receives input and other requests 115 from a plurality of customer organizations 105A-C via network 125 (such as a public Internet). For example, incoming database queries, API requests, interactions with displayed graphical user interfaces and displays at the client devices 106A-C, or other inputs may be received from the customer organizations 105A-C to be processed against the multi-tenant database system 130, including queries to read, write, and update data stored within the databases 155.

In one embodiment, each customer organization 105A-C is an entity selected from the group consisting of: a separate and distinct remote organization, an organizational group within the host organization 110, a business partner of the host organization 110, or a customer organization 105A-C that subscribes to cloud computing services provided by the host organization 110.

In one embodiment, requests 115 are received at, or submitted to, a web-server 175 within host organization 110. Host organization 110 may receive a variety of requests for processing by the host organization 110 and its multi-tenant database system 130. Incoming requests 115 received at web-server 175 may specify which services from the host organization 110 are to be provided, such as query requests, search request, status requests, database transactions, graphical user interface requests and interactions, processing requests to retrieve, update, or store data on behalf of one of the customer organizations 105A-C, code execution requests, and so forth. Web-server 175 may be responsible for receiving requests 115 from various customer organizations 105A-C via network 125 and provide a web-based interface or other graphical displays to an end-user client device 106A-C or machine originating such data requests 115.

Authenticator 140 operates on behalf of the host organization to verify, authenticate, and otherwise credential users attempting to gain access to the host organization.

Query interface 180 provides functionality to pass queries from web-server 175 into the multi-tenant database system 130 for execution against the databases 155 or other data stores of the host organization's production environment 111. In one embodiment, the query interface 180 implements an Application Programming Interface (API) through which queries may be executed against the databases 155 or other data stores. Query optimizer 160 performs query translation and optimization, for instance, on behalf of other functionality such as functionality of a graphical interface which possesses sufficient information to architect a query yet lacks the necessary logic to actually construct the appropriate query syntax. In other instances, query optimizer 160 modifies a submitted query to optimize its execution within the host organization without affecting the resulting dataset returned responsive to such an optimized query. Analysis engine 185 operates on behalf of functionality within the host organization to perform additional data evaluation where necessary.

Host organization 110 may implement a request interface 176 via web-server 175 or as a stand-alone interface to receive requests packets or other requests 115 from the client devices 106A-C. Request interface 176 further supports the return of response packets or other replies and responses 116 in an outgoing direction from host organization 110 to the client devices 106A-C.

Historical record viewer 186 and corruption user interface (corruption UI) 187 are further depicted as interacting with the databases 155. The historical record viewer 186 enables users to view past transactions at the databases as well as the past or historical state of data reflected by their database records. Corruption UI 187 enables users to select past transactions to the databases to revert or modify, as well as view what impact upon other records their proposed changes may cause and then to confirm compensating database transactions to correct logical corruption as will be described in additional detail below. Changes available for view by any particular database tenant or user are restricted to their view of the data in the database as controlled by conventional access restrictions, ownership, and responsibility for such data. Thus, one tenant cannot see the data of another tenant unless they are given access rights to view such data. Thus, according to certain embodiments, both OrgID and UserID are included within the transactions logs to correctly align access rights to the view granted to any user via the historical record viewer 186 and the corruption user interface 187.

FIG. 2 depicts another exemplary architecture 200 in accordance with described embodiments. Depicted here the execution hardware, software, and logic 120 of the multi-tenant database system 130, but additionally depicted are LSM databases 255 operating as data stores for the multi-tenant database system 130. The historical record viewer 186 is again depicted which interacts with the logs 257 of the LSM database 255 and further interacts with the corruption user interface 187.

According to the described embodiments, the LSM databases 255 contain stored records 256 and operate utilizing LSM database logs 257 and immutable LSM trees containing database transactions that are intended to be merged into the stored records 256 of the LSM databases 255. Database records are stored within database blocks. Database blocks in turn are contained within database extents. Extents are groups of blocks that contain groups of rows. The records are the changes to the database including inserted rows, updated rows and deleted rows.

LSM databases 255 queue up these pending transactions rather than immediately writing them into their permanent location in the database and subsequent processing merges the extents 258 into the stored records 256 of the LSM databases 255, for instance, by collapsing updates and modifications to rows, effecting new row insertions, and permanently effecting row deletions. For instance, as depicted, a record update 223 is shown being written to the extent 258 of the LSM databases 255 and a subsequent merge 224 causes the record update 223 to be effected permanently into the stored records 256 of the LSM database 255, likely with many other pending/queued transactions, be they insertions, modifications, or deletions.

Many conventional databases operate utilizing an update-in-place methodology in which any changes to the database requires the database engine to read out an affected record, update the record, and write the updated record back into the database at the correct location. With an update-in-place database the database blocks are not immutable, that is, they can be changed. Tables are a logical construct on top of the database blocks but the records have to be written to database blocks thus changing an existing block to a new version within an update-in-place database. Records and rows are written to a block and will typically remain in that block until changed or until the update-in-place database is reorganized. LSMs databases on the other hand defer placement of the records and subsequently move records as part

of merge processing. Record inserts to such update-in-place databases must be written into the correct location in the correct block at the time of insertion, changing the version of the block, and similarly, record deletions must remove the correct record from the correct block at the time of deletion, again changing the version of the block.

Logging exists for such conventional update-in-place database systems with enough information to repeat the entire operation if the direct write to the block gets interrupted mid-way through (e.g., by a system crash) via an operation called “replaying” the log, thus requiring that each transaction in an update-in-place database be written to disk twice: once to the log, and once into the permanent location within the correct database block of the update-in-place database. Notwithstanding the logging by update-in-place databases, the transactions are effected into the database at the exact and permanent location at the time that the transaction occurs.

If a block becomes corrupted in such an update-in-place database the entire database will often crash to prevent the corruption from spreading, although as noted above, some enterprise level databases implement database block isolation schemes to prevent crashing. Such enterprise level databases can sustain corruptions and even self-repair themselves under certain circumstances, although such technology is extremely expensive to implement. Oracle™ for example requires a second such enterprise system with a fully redundant set of data that is synchronized and available at all times such that the redundant set can be used to repair corruption. Certain mainframe databases also will sustain corruption without crashing such as those systems support financial transactions, but again, the technology is extremely expensive to implement. For instance, banks and stock markets operate such mainframe databases.

An interesting case may occur in an update-in-place database where a database block is damaged then because of the damage, it is unknown which of the thousands of tenants having data stored on the database may be affected by such damage. Such a problem is especially true for an update-in-place database where the tenants are inter-mixed, potentially without ordering. Consider therefore performing a scanning query for “where tenant=<tenantID>” and encountering a corrupted block. Without knowing which tenants are in the block, the scan must fail to avoid returning an incorrect SQL statement result which could then propagate logical corruption throughout the update-in-place database.

Conversely, with an LSM database where the extents and blocks have ordered blocks with ordered tenant data within them it is at least knowable which tenants may be associated with and thus affected by a corrupted block and thus, it can then be decided whether or not the scan should continue or error out based on whether the particular tenant running the scan is one of those that would be affected by the corruption to that block. Even if the block is corrupted or entirely lost it is still possible to determine whether a tenant executing the scan will be affected by the corruption because the scan logic can look to immediately adjacent blocks within the extent to see the ordering of tenants and the range of tenants could potentially be affected by the lost or corrupted block. In such a situation, scanning can therefore continue without error where it is determined that a tenant executing the scan does not fall within the range of potentially affected tenants, or error out when appropriate should the tenant unfortunately fall within the range. If tenants do fall within the range then the corruption can at least be identified and trigger corruption rectification procedures be that a replay, compensating transaction, extent replacement, etc.

In accordance with certain embodiments therefore, the LSM database orders records in blocks by tenant and within extents. In such an embodiment, scans will naturally skip over blocks that are not needed, made possible by the ordering which not only greatly improves scan performance but additionally permits tenants having no association with a block to benefit from scans that continue successfully scanning as long as they are orthogonal to the corrupted block data, notwithstanding the presence of corrupted blocks in the ordered LSM database. In such a way, the LSM database continues to operate in a live production state capable of transacting on behalf of tenants whose blocks are not afflicted with such corruption.

The ability for the LSM database to continue in a live operational state notwithstanding the presence of corrupted blocks is critical to the nature of the multi-tenant database system environment which services tends of thousands of diverse customers concurrently. Consider for example a massive database within the multi-tenant database system having an exemplary 30 TB of data hosted on behalf of 1000 tenants. The particular corruption issue may reside within a block that affects just a single one of the 1000 tenants or perhaps a dozen or so of the tenants, yet with conventional update-in-place databases, the corrupted block will trigger the entire database to crash (assuming it is not a corruption resilient enterprise database) thus taking all 1000 tenants offline for a problem that approximately ~99% of the tenants do not care about and for which such tenants may never be directly affected by the corruption because they are in no way affiliated with the data block found to be corrupted.

Consider further a database supporting millions of financial customers simultaneously, such as within a banking and finance environments. Crashing the database for corruption at a single block would be flatly unacceptable as doing so could cause disastrous financial repercussions for the bank or even the larger economy depending on the scope and nature of the financial transactions supported by the system. For instance, consider the havoc should Visa and MasterCard’s entire transaction processing system crash or worse yet, damage to financial markets should a stock exchange crash. Therefore, an update-in-place database in such an environment may maintain the system’s availability rather than crashing upon the detection of a corrupted block and instead attempt to isolate the corruption from non-affected customers during the database’s operational runtime so that the single corrupted block, and the few customers affected, can be remedied without triggering a full-blown service outage.

Unfortunately, resolving corruption in such a situation is enormously complex due to the nature of the update-in-place database’s propensity to be constantly changing the blocks within which records are stored as the database transactions are processed. While the update-in-place database continues operating, transactions are constantly updating the permanent location of data in the database this making it very risky and technically complex to rectify corrupted blocks without taking the database offline.

In order to fix the corruption it necessary to obtain or reconstruct the correct version of a database block. Reconstructing the incorrect version will introduce difficult to detect logical corruption which is in of itself a disastrous consequence to the data integrity of the database, yet to make matters worse, the logical corruption resulting from reconstructing the wrong version of the corrupted block will then logically propagate across the database in a manner that is extremely difficult to predict and potentially impossible to subsequently detect rendering the full impact of the corrup-

tion potentially unknowable. For instance, block reads are not logged or tracked and as such, there would be no record of those transactions having pulled data from the logically corrupted block, used that data in some way within an application layer separate from the database, and then written data back into the database having been indirectly affected and corrupted by the erroneous block version reconstruction. For instance, if perhaps a person's salary record is corrupted, and then a manual update to fix the corruption enters incorrect data, the erroneous block version update will logically corrupt that block but render it physically valid, appearing as though the corruption is fixed, yet could unknowingly give somebody a raise or reduction in salary, or remove their salary completely, or allocate their salary to another person, etc.

For these reasons, correcting corruption on a live operational update-in-place database, while possible, is a difficult and dangerous proposition that typically requires very skilled and expensive experts as well as support from the database vendor all while risking the introduction of new logical corruption into the database.

LSM database 255 technology is fundamentally different than update-in-place database technology. A LSM or LSM-tree type database is a "Log-Structured Merge-tree" type database and does not write updates, deletions, or insertions directly into the permanent database location at the time of the transaction as is done with an update-in-place database. Instead of writing the database transactions into the permanent location in the database and writing the entry into a log also as is done with an update-in-place database, an LSM database 255 simply writes the database transaction (e.g., such as record update 223) to the end of a queue, referred to as an extent, and is then finished with processing for that given database transaction for the time being. The LSM database additionally writes checkpoints to disk periodically which describe the state of the file system at that point in time providing organization to the queued transactions stored in the extents 258 so as to make the data associated with pending updates more efficient to find. An extent 258 is one or more blocks of data ("blocks" are sometimes also referred to as pages or fragments). If the LSM database 255 were to use only blocks then there would be millions of blocks to track and data would become increasingly difficult to find with organizational overhead costs increasing dramatically. By placing blocks into extents 258, each having, for example, 300 to 500 blocks, then the computational performance of organization and indexing to find data is improved dramatically. This is necessary because notwithstanding the fact that record updates 223 and other transactions are written to the extents 258 before subsequent merge operations, follow-on queries to the stored records 256 of the LSM database 255 must accurately reflect the state of the data stored within them, regardless of whether the latest version of the record exists as a merged and thus permanent stored record 256 or exists as a pre-merge queued database transaction within one of the extents 258.

The LSM database system engine can be thought of as a lazy accountant. Rather than locating the exact row in a table, reading it out, updating it, and writing it back into the proper and permanent location at the time of the transaction, as is required by the update-in-place database, the LSM database's lazy accountant instead logs and queues all the transactions with the intention of getting around to them later. This results in many small "extents" 258 which include database transactions (e.g., such as the record update 223) that need to subsequently be inserted or updated into the

LSM database's permanent location at a later time, through a process called merging 224.

Therefore, if a transaction for an LSM database deletes a record from the database, the deletion will be logged and immutably recorded into the extent 258 at the time of the transaction, however, the record still remains at its permanent location within the stored records 256 of the LSM database 255, until such time that the extent and its transactions are processed or merged 224 into the LSM databases 255 at which point in time the record deletion is made effective or permanent within the stored records 256 of the LSM database 255. Once made effective or permanent through merging 224, the extent 258 is no longer required and is therefore deleted or marked for deletion. However, transactions may nevertheless be logged or the extents themselves may be kept as logs for redundancy and future reference, for example, for use by the historical record viewer 186 and corruption user interface 187.

FIG. 3 depicts another exemplary architecture 300 in accordance with described embodiments. As before, there remains the execution hardware, software, and logic 120 of the multi-tenant database system 130 and the LSM databases 255, its stored records 256, logs 257, and extents 258. Additionally depicted is the corruption editor 330 which interacts with the corruption user interfaced 187 and the historical record viewer 186.

With the LSM database 255, any changes subsequent to a checkpoint can be recovered by replaying a relatively small number of log 257 entries following the checkpoint. Sufficiently large logs 257 permit a view into the past to see what transactions occurred and thus enabling those past transactions to be replayed or undone via compensating transactions 334. Larger logs 257 thus equate to a longer available history of transactions that may be referenced, viewed, or displayed to a user via the historical record viewer 186. While it may be infeasible to store all transactions in the logs 257 indefinitely, it is practical to store several days worth of transactional logs, thus permitting replay or review and correction for any transaction during that past period of time as reflected by the available transaction logs. In such a way the transaction logs provide for data redundancy within the system enabling a variety of mechanisms to fix corruption in an LSM type database, as described herein.

As depicted here, the corruption editor 330 enables a user (e.g., via the corruption UI 187 and historical record viewer 186) to view and retrieve compensating transaction data 333 from either the stored records 256 or the logs 257, or both, and then via the corruption editor 330 a compensating transaction 334 can be generated and issued to the LSM databases 255 to correct or counteract corruption.

As noted previously, the queued up database transactions of an LSM database 255 are kept within the immutable extents 258 representing the logged transactions that are intended to be merged into the LSM database but have not yet been handled. The extent 258 cannot ever be modified. Instead, it exists until the LSM database merges the queued transactions at which point it no longer required or relevant and thus deleted or marked for deletion depending on how the log 257 handles the processed extents 258.

Thus, unlike the update-in-place database which is in a constant state of change due to the updates being transacted to the physical locations of the database and in which corruption can quickly spread, the LSM database 255, if corrupted, will localize that corruption to a particular extent 258 yet to be merged, at least for a period of time. Resolving corruption within the LSM database 255 while it remains live and operational is therefore much more feasible than

with the update-in-place database. For instance, once corruption is identified, it can be resolved by replacing the pending extent **258** with a corrected or replacement extent **335** having the corrupted data re-stamped, or re-applied in its correct form. The corrected/replacement extent **335** may be re-generated by replaying the necessary transactions as recorded by the logs **257** and updating the corrected/replacement extent **335** based on the compensating transaction data **333**, after which the corruption editor **330** causes the corrected/replacement extent **335** to over-write the pending extent **258** prior to merge. Alternatively, compensating transactions **334** can be queued into the latest extent **258** without having to replace any existing pre-merge extent, in which case the compensating transactions **334** which would simply counter-act the corrupting transaction.

With sufficient redundancy in the logs and/or older extents stored after merge processing, including extents having since been merge and marked for deletion, it is possible to replay the logs again creating the same extent that was created previously based on compensating transaction data **333** via a “replay” or to generate new compensating transactions **334** which will in turn cause a compensating merge event through the LSM databases’ **255** normal processing thus re-creating, restoring, or otherwise correcting the corruption issue. In the event that records were erroneously or maliciously inserted a compensating transaction **334** would simply delete the new record and at merge processing the new record would be effectively destroyed or negated by the compensating transaction **334**.

Conversely, if a record deletion is transacted against the database, but is in error and thus represents logical corruption, the enqueued record deletion cannot be removed from the immutable extent’s **258** pending merge processing. However, the erroneously deleted record still exists within the LSM database’s **255** stored records **256** at its permanent location until such time that the merge occurs for the extent **258** having the erroneous delete transaction queued therein. Therefore, a compensating transaction **334** can read out that record from the stored records **256** as compensating transaction data **333** and then transact it for entry into the database via the compensating transaction **334**, thus queuing it for later merge via the extent **258**. When the LSM database **255** performs its subsequent merge, the two transactions will effectively annihilate one another when fully processed. More particularly, the first erroneous delete will be effected into the LSM database **255** via the merge and then the insertion will be effected into the LSM database **255** via the merge, resulting in a net no-change event from the perspective of a user or anyone querying for the data belonging to the erroneously deleted record.

The compensating transactions **334** or corrected/replacement extent **335** operations may require special handling, bypassing the standard SQL interface available to database users, however, the fact that transactions are not immediately updated at their permanent stored records **256** location in the LSM database **255** makes such special handling during live operation of the LSM database **255** feasible where it is not always possible with an update-in-place database. For instance, special queries can be executed against the LSM database **255** to query the permanent stored records **256** location, bypassing the logged/enqueued modification present in an extent **258**, so as to obtain the compensating transaction data **333** of a record pending deletion. One way of doing this differentiate between an older record pending deletion and the transaction specifying the deletion based on time-stamps kept by the LSM database **255**. A normal insert transaction can then be issued as a

compensating transaction **334**. Alternatively, while the LSM’s extent **258** having the queued transactions is itself immutable and cannot be modified, where corruption exists in such an extent **258**, a corrected/replacement extent **335** can be substituted reflecting the corruption fix, which will then be taken up by the LSM database and merged in the usual fashion. Again, this may require special handling transactions, but can be implemented through customized tools and interfaces to carry out the requisite operations such as the corruption editor **330** and its supporting historical record viewer **186** and corruption user interface **187**.

FIG. 4 depicts an exemplary graphical interface operating at a mobile, smartphone, or tablet computing device in accordance with the embodiments. In particular, there is depicted a smartphone or tablet computing device **401** having embodied therein a touch interface **405**, such as a mobile display. Presented or depicted to the mobile display **405** is the historical record viewer **402** in which the various changed records **425** are depicted to a user via the display and additionally presented or depicted to the mobile display **405** is the corruption user interface **480** with a record editor **484** from which a user can select one of the changed records **425** from the historical record viewer **402** and then modify the selected changed record at the corruption user interface **480** via the record editor **484**. For instance, the user can, via the mobile display **405**, change the values of the selected changed record at the record editor **484**, restore the record via the undo delete button, delete the record via the undo insert button, commit edits entered to the record editor **484**, cancel, or select from various other options appropriate for the UI. At element **481** the user can also select the option to preview how the proposed changes impact other records.

The LSM database **255** stores logical information about where data is stored whereas an update-in-place database conversely stores physical information about where data is stored. Therefore, embodiments further benefit from the LSM’s structure by permitting users to undo or compensate for erroneous transactions or even malicious transactions carried out by, for example, a disgruntled employee or a hacker having gained illicit access to the user’s account or the customer organization’s data. For example, a disgruntled employee or hacker may go into a database and delete records or modify data, such as salary information, etc. These changes are perfectly valid in a physical sense. That is to say, they are transacted to the database and written to the proper location regardless of whether the database is an update-in-place or LSM database. However, such changes represent logical errors, which can be corrected as described herein by reviewing or replaying available logs to show what changes were made via historical record viewer **402** and corruption UI **480** and then permitting the user to undo those changes through compensating transactions **334** via the functionality of the corruption editor **330** as set forth at FIG. 3.

Consider another example in which a user’s account was accessed at 3:00 AM by a hacker with a variety of database transactions issued against the LSM database from the user’s account. As noted previously, such transactions are perfectly valid from a physical sense, but may very well represent logical corruption in the database due to the fact that the transactions are unwanted and erroneously applied to the database. In certain embodiments the user is notified of transactions out of the ordinary, such as edits being made to the customer organization’s data at an odd time, such as 3:00 AM, or from a strong location or IP address. These and other

means may bring to the user's attention the fact that some logical corruption may have been introduced into the database.

Using the logs, the historical record viewer **402** presents user a user with a display showing all transactions made within a period of time. For instance, via the interfaces presented at the mobile display **405** the user can select changed records **425** ranging in time from, for example, 3:00 AM to 6:00 AM, causing the display and UI logic to retrieve those transactions from the log and display them to the user for review or editing via record editor **484**. The LSM database itself possesses no knowledge that such properly executed transactions represent logical corruption, however, the user, customer organization, or tenant having ownership or responsibility for such data would have intimate knowledge of the data. Such persons can retrieve the data and review it for correctness. If such data is incorrect, then according to described embodiments, the user can select via the UI which records to revert, for instance, by checking check boxes, or highlighting records, clicking the appropriate buttons, etc. The UI will then institute the proper compensating transactions on behalf of the user (e.g., reading the records at a time prior to the 3:00 AM hack and re-updating them or re-inserting them with the proper data record, etc) to revert the LSM database records to their prior state or to an updated state at the discretion of the user.

According to another embodiment, the user can additionally be presented with a view at the UI of all records that will be affected by the proposed changes by checking the selection at element **481**. For instance, a user after selecting which records to revert or modify can be presented with a verification or confirmation screen at the UI asking them to confirm the proposed changes and showing the user those records that will be modified or directly affected by the compensating transactions. For instance, it may be that subsequent database transactions were processed subsequent to a malicious hack, and the proposed changes, while reverting the affects of the erroneous transactions, will also revert properly processed transactions subsequent to the hack, and may therefore require additional review or modification of the proposed compensating transactions by the user, all of which is enabled at the UI interface presented to the user in accordance with such embodiments.

Physical corruption issues or logical corruption may also occur at the database engine level. For instance, a record may be created or deleted without a corresponding change to the index. Therefore, in accordance with another embodiment, a count or checksum is maintained for rows inserted into the tables of the database and a corresponding count or checksum is maintained for rows of the tables in the database represented by an index. A check may then verify whether the two are in sync, and if not, then a user may be notified to review transactions corresponding in time when the counts or checksums became out of sync.

CRC checks may additionally be maintained on blocks to check for corruption and then if corruption is detected through a scan, users affected by the corrupted block can be notified via an error message or other trigger without requiring the entire database be taken offline and impacting those users having no data associated the corrupted block. In such a way, only a limited number of rows, records, or transactions that are actually impacted by the corruption need to be isolated thus leaving the system available for the overwhelming majority of other records and users.

FIG. 5A depicts a tablet computing device **501** and a hand-held smartphone **502** each having a circuitry integrated therein as described in accordance with the embodiments.

As depicted, each of the tablet computing device **501** and the hand-held smartphone **502** include a touch interface **503** (e.g., a touchscreen or touch sensitive display) and an integrated processor **504** in accordance with disclosed embodiments.

For example, in one embodiment, a system embodies a tablet computing device **501** or a hand-held smartphone **502**, in which a display unit of the system includes a touchscreen interface **503** for the tablet or the smartphone and further in which memory and an integrated circuit operating as an integrated processor are incorporated into the tablet or smartphone, in which the integrated processor implements one or more of the embodiments described herein. In one embodiment, the integrated circuit described above or the depicted integrated processor of the tablet or smartphone is an integrated silicon processor functioning as a central processing unit (CPU) and/or a Graphics Processing Unit (GPU) for a tablet computing device or a smartphone.

FIG. 5B is a block diagram **500** of an embodiment of tablet computing device, a smart phone, or other mobile device in which touchscreen interface connectors are used. Processor **510** performs the primary processing operations. Audio subsystem **520** represents hardware (e.g., audio hardware and audio circuits) and software (e.g., drivers, codecs) components associated with providing audio functions to the computing device. In one embodiment, a user interacts with the tablet computing device or smart phone by providing audio commands that are received and processed by processor **510**.

Display subsystem **530** represents hardware (e.g., display devices) and software (e.g., drivers) components that provide a visual and/or tactile display for a user to interact with the tablet computing device or smart phone. Display subsystem **530** includes display interface **532**, which includes the particular screen or hardware device used to provide a display to a user. In one embodiment, display subsystem **530** includes a touchscreen device that provides both output and input to a user.

I/O controller **540** represents hardware devices and software components related to interaction with a user. I/O controller **540** can operate to manage hardware that is part of audio subsystem **520** and/or display subsystem **530**. Additionally, I/O controller **540** illustrates a connection point for additional devices that connect to the tablet computing device or smart phone through which a user might interact. In one embodiment, I/O controller **540** manages devices such as accelerometers, cameras, light sensors or other environmental sensors, or other hardware that can be included in the tablet computing device or smart phone. The input can be part of direct user interaction, as well as providing environmental input to the tablet computing device or smart phone.

In one embodiment, the tablet computing device or smart phone includes power management **550** that manages battery power usage, charging of the battery, and features related to power saving operation. Memory subsystem **560** includes memory devices for storing information in the tablet computing device or smart phone. Connectivity **570** includes hardware devices (e.g., wireless and/or wired connectors and communication hardware) and software components (e.g., drivers, protocol stacks) to the tablet computing device or smart phone to communicate with external devices. Cellular connectivity **572** may include, for example, wireless carriers such as GSM (global system for mobile communications), CDMA (code division multiple access), TDM (time division multiplexing), or other cellular service standards). Wireless connectivity **574** may include,

for example, activity that is not cellular, such as personal area networks (e.g., Bluetooth), local area networks (e.g., WiFi), and/or wide area networks (e.g., WiMax), or other wireless communication.

Peripheral connections **580** include hardware interfaces and connectors, as well as software components (e.g., drivers, protocol stacks) to make peripheral connections as a peripheral device (“to” **582**) to other computing devices, as well as have peripheral devices (“from” **584**) connected to the tablet computing device or smart phone, including, for example, a “docking” connector to connect with other computing devices. Peripheral connections **580** include common or standards-based connectors, such as a Universal Serial Bus (USB) connector, DisplayPort including MiniDisplayPort (MDP), High Definition Multimedia Interface (HDMI), Firewire, etc.

FIG. 6 is a flow diagram illustrating a method **600** for fixing logical or physical corruption in databases using immutable LSM trees in accordance with disclosed embodiments. Method **600** may be performed by processing logic that may include hardware (e.g., circuitry, dedicated logic, programmable logic, microcode, etc.), software (e.g., instructions run on a processing device to perform various operations such as storing records, processing transactions, recording, logging, identifying errors and corruption, correcting errors and corruption, displaying, retrieving, communicating, updating, applying, committing, querying, executing, generating, exposing, transmitting, sending, returning, etc., in pursuance of the systems, apparatuses, and methods, as described herein. For example, the computing architecture (e.g., within production environment **111**) of host organization **110** as depicted at FIG. 1, the hand-held smartphone **502** or mobile tablet computing device **501** depicted at FIG. 5A, the machine **800** at FIG. 8, or the system **900** at FIG. 9, may implement the described methodologies. Some of the blocks and/or operations listed below are optional in accordance with certain embodiments. The numbering of the blocks presented is for the sake of clarity and is not intended to prescribe an order of operations in which the various blocks must occur.

At block **605**, processing logic stores records in a LSM database at a host organization.

At block **610**, processing logic processes transactions to the LSM database by writing the transactions into immutable extents for subsequent merge into the LSM database.

At block **615**, processing logic records a history of the transactions in a log.

At block **620**, processing logic identifies corruption in at least one of the records of the LSM database represented within the immutable extents prior to merge of the immutable extents to the LSM database.

At block **625**, processing logic corrects the identified corruption to the at least one of the records of the LSM database based at least in part on the recorded history of the transactions in the log during live operation of the LSM database.

According to another embodiment of method **600**, correcting the identified corruption to the at least one of the records of the LSM database includes at least one of: restoring a modified record to a prior state; deleting an erroneously inserted record; restoring an erroneously deleted record; or changing one or more field values of an erroneously modified record.

According to another embodiment of method **600**, correcting the identified corruption to the at least one of the records of the LSM database includes: reading a stored record from the LSM database having a corresponding

transaction pending merge processing within the immutable extents; reading compensating transaction data from the log; generating a compensating transaction based on the compensating transaction data from the log and the stored record from the LSM database; and sending the compensating transaction to the LSM database.

According to another embodiment of method **600**, sending the compensating transaction to the LSM database includes: issuing the compensating transaction to the LSM database as a new transaction, in which the new transaction is enqueued within the immutable extent of the LSM database for subsequent merge processing; and in which the subsequent merge processing counteracts the corresponding transaction pending merge processing within the immutable extents for the stored record read from the LSM database.

According to another embodiment of method **600**, sending the compensating transaction to the LSM database includes: generating a replacement extent by replaying a plurality of transactions from the log for one of the immutable extents of the LSM database awaiting subsequent merge and updating the identified corruption in the replacement extent using the compensating transaction; and replacing the one immutable extent for which the plurality of transactions were replayed with the replacement extent.

According to another embodiment of method **600**, identifying corruption in at least one of the records of the LSM database includes: triggering error messages to users for any queries to the LSM database having the identified corruption within the query claim scope; and maintaining availability and accessibility to the LSM database to database queries.

According to another embodiment, method **600** further includes: presenting a historical records view to a user interface based on the log, the historical records view displaying one or more records updated during a time range configurable at the user interface.

According to another embodiment, method **600** further includes: receiving user input selecting one or more of the records updated during the time range; and receiving user input specifying corruption edits to the one or more records.

According to another embodiment of method **600**, the user input specifying corruption edits to the one or more records includes user input specifying one or more of: a selected record erroneously inserted into the LSM database to be deleted from the LSM database via a compensating delete transaction; a selected record erroneously deleted from the LSM database to be restored to the LSM database via a compensating insert transaction; a selected record erroneously modified within the LSM database to be reverted to a prior state via a compensating update transaction; and a selected record erroneously modified within the LSM database to be updated to a new state via a compensating update transaction.

According to another embodiment, method **600** further includes: presenting a preview of other records impacted by a proposed compensating transaction to the user interface prior to proceeding with the proposed compensating transaction.

According to another embodiment, method **600** further includes: receiving a selection from the user interface specifying one of the displayed records updated during the time range; presenting a record editor to the user interface having the record corresponding to the selection populated therein; receiving user input to the fields of the record corresponding to the selection from the record editor at the user interface; and generating a compensating transaction for the record corresponding to the selection using the user input to the fields of the record.

According to another embodiment of method 600, the host organization implements the method via computing architecture of the host organization including at least the processor and the memory; in which a user interface operates at a user client device remote from the host organization and communicatively interfaces with the host organization via a public Internet; and in which the host organization operates as a cloud based service provider to the user client device.

According to another embodiment of method 600, the host organization provides a multi-tenant database system via the LSM database and the computing architecture of the host organization, the multi-tenant database system having elements of hardware and software that are shared by a plurality of separate and distinct customer organizations, each of the separate and distinct customer organizations being remotely located from the host organization.

According to another embodiment of method 600, identifying corruption in at least one of the records of the LSM database includes one of: identifying corruption responsive to a physical corruption error triggered by the LSM database; or identifying corruption responsive user input at a corruption user interface, the user input indicating a selected record is to be corrected via a compensating transaction.

According to a particular embodiment, there is non-transitory computer readable storage media having instructions stored thereon that, when executed by a processor in a host organization, the instructions cause the host organization to perform operations including: storing records in a LSM database at the host organization; processing transactions to the LSM database by writing the transactions into immutable extents for subsequent merge into the LSM database; recording a history of the transactions in a log; identifying corruption in at least one of the records of the LSM database represented within the immutable extents prior to merge of the immutable extents to the LSM database; and correcting the identified corruption to the at least one of the records of the LSM database based at least in part on the recorded history of the transactions in the log during live operation of the LSM database.

FIG. 7A illustrates a block diagram of an environment 798 in which an on-demand database service may operate in accordance with the described embodiments. Environment 798 may include user systems 712, network 714, system 716, processor system 717, application platform 718, network interface 720, tenant data storage 722, system data storage 724, program code 726, and process space 728. In other embodiments, environment 798 may not have all of the components listed and/or may have other elements instead of, or in addition to, those listed above.

Environment 798 is an environment in which an on-demand database service exists. User system 712 may be any machine or system that is used by a user to access a database user system. For example, any of user systems 712 can be a handheld computing device, a mobile phone, a laptop computer, a work station, and/or a network of computing devices. As illustrated in FIG. 7A (and in more detail in FIG. 7B) user systems 712 might interact via a network 714 with an on-demand database service, which is system 716.

An on-demand database service, such as system 716, is a database system that is made available to outside users that do not need to necessarily be concerned with building and/or maintaining the database system, but instead may be available for their use when the users need the database system (e.g., on the demand of the users). Some on-demand database services may store information from one or more

tenants stored into tables of a common database image to form a multi-tenant database system (MTS). Accordingly, “on-demand database service 716” and “system 716” is used interchangeably herein. A database image may include one or more database objects. A relational database management system (RDMS) or the equivalent may execute storage and retrieval of information against the database object(s). Application platform 718 may be a framework that allows the applications of system 716 to run, such as the hardware and/or software, e.g., the operating system. In an embodiment, on-demand database service 716 may include an application platform 718 that enables creation, managing and executing one or more applications developed by the provider of the on-demand database service, users accessing the on-demand database service via user systems 712, or third party application developers accessing the on-demand database service via user systems 712.

The users of user systems 712 may differ in their respective capacities, and the capacity of a particular user system 712 might be entirely determined by permissions (permission levels) for the current user. For example, where a salesperson is using a particular user system 712 to interact with system 716, that user system has the capacities allotted to that salesperson. However, while an administrator is using that user system to interact with system 716, that user system has the capacities allotted to that administrator. In systems with a hierarchical role model, users at one permission level may have access to applications, data, and database information accessible by a lower permission level user, but may not have access to certain applications, database information, and data accessible by a user at a higher permission level. Thus, different users will have different capabilities with regard to accessing and modifying application and database information, depending on a user’s security or permission level.

Network 714 is any network or combination of networks of devices that communicate with one another. For example, network 714 can be any one or any combination of a LAN (local area network), WAN (wide area network), telephone network, wireless network, point-to-point network, star network, token ring network, hub network, or other appropriate configuration. As the most common type of computer network in current use is a TCP/IP (Transfer Control Protocol and Internet Protocol) network, such as the global internet network of networks often referred to as the “Internet” with a capital “I,” that network will be used in many of the examples herein. However, it is understood that the networks that the claimed embodiments may utilize are not so limited, although TCP/IP is a frequently implemented protocol.

User systems 712 might communicate with system 716 using TCP/IP and, at a higher network level, use other common Internet protocols to communicate, such as HTTP, FTP, AFS, WAP, etc. In an example where HTTP is used, user system 712 might include an HTTP client commonly referred to as a “browser” for sending and receiving HTTP messages to and from an HTTP server at system 716. Such an HTTP server might be implemented as the sole network interface between system 716 and network 714, but other techniques might be used as well or instead. In some implementations, the interface between system 716 and network 714 includes load sharing functionality, such as round-robin HTTP request distributors to balance loads and distribute incoming HTTP requests evenly over a plurality of servers. At least as for the users that are accessing that

server, each of the plurality of servers has access to the MTS' data; however, other alternative configurations may be used instead.

In one embodiment, system 716, shown in FIG. 7A, implements a web-based customer relationship management (CRM) system. For example, in one embodiment, system 716 includes application servers configured to implement and execute CRM software applications as well as provide related data, code, forms, webpages and other information to and from user systems 712 and to store to, and retrieve from, a database system related data, objects, and Webpage content. With a multi-tenant system, data for multiple tenants may be stored in the same physical database object, however, tenant data typically is arranged so that data of one tenant is kept logically separate from that of other tenants so that one tenant does not have access to another tenant's data, unless such data is expressly shared. In certain embodiments, system 716 implements applications other than, or in addition to, a CRM application. For example, system 716 may provide tenant access to multiple hosted (standard and custom) applications, including a CRM application. User (or third party developer) applications, which may or may not include CRM, may be supported by the application platform 718, which manages creation, storage of the applications into one or more database objects and executing of the applications in a virtual machine in the process space of the system 716.

One arrangement for elements of system 716 is shown in FIG. 7A, including a network interface 720, application platform 718, tenant data storage 722 for tenant data 723, system data storage 724 for system data 725 accessible to system 716 and possibly multiple tenants, program code 726 for implementing various functions of system 716, and a process space 728 for executing MTS system processes and tenant-specific processes, such as running applications as part of an application hosting service. Additional processes that may execute on system 716 include database indexing processes.

Several elements in the system shown in FIG. 7A include conventional, well-known elements that are explained only briefly here. For example, each user system 712 may include a desktop personal computer, workstation, laptop, PDA, cell phone, or any wireless access protocol (WAP) enabled device or any other computing device capable of interfacing directly or indirectly to the Internet or other network connection. User system 712 typically runs an HTTP client, e.g., a browsing program, such as Microsoft's Internet Explorer browser, a Mozilla or Firefox browser, an Opera, or a WAP-enabled browser in the case of a smartphone, tablet, PDA or other wireless device, or the like, allowing a user (e.g., subscriber of the multi-tenant database system) of user system 712 to access, process and view information, pages and applications available to it from system 716 over network 714. Each user system 712 also typically includes one or more user interface devices, such as a keyboard, a mouse, trackball, touch pad, touch screen, pen or the like, for interacting with a graphical user interface (GUI) provided by the browser on a display (e.g., a monitor screen, LCD display, etc.) in conjunction with pages, forms, applications and other information provided by system 716 or other systems or servers. For example, the user interface device can be used to access data and applications hosted by system 716, and to perform searches on stored data, and otherwise allow a user to interact with various GUI pages that may be presented to a user. As discussed above, embodiments are suitable for use with the Internet, which refers to a specific global internetwork of networks. However, it is understood

that other networks can be used instead of the Internet, such as an intranet, an extranet, a virtual private network (VPN), a non-TCP/IP based network, any LAN or WAN or the like.

According to one embodiment, each user system 712 and all of its components are operator configurable using applications, such as a browser, including computer code run using a central processing unit such as an Intel Pentium® processor or the like. Similarly, system 716 (and additional instances of an MTS, where more than one is present) and all of their components might be operator configurable using application(s) including computer code to run using a central processing unit such as processor system 717, which may include an Intel Pentium® processor or the like, and/or multiple processor units.

According to one embodiment, each system 716 is configured to provide webpages, forms, applications, data and media content to user (client) systems 712 to support the access by user systems 712 as tenants of system 716. As such, system 716 provides security mechanisms to keep each tenant's data separate unless the data is shared. If more than one MTS is used, they may be located in close proximity to one another (e.g., in a server farm located in a single building or campus), or they may be distributed at locations remote from one another (e.g., one or more servers located in city A and one or more servers located in city B). As used herein, each MTS may include one or more logically and/or physically connected servers distributed locally or across one or more geographic locations. Additionally, the term "server" is meant to include a computer system, including processing hardware and process space(s), and an associated storage system and database application (e.g., OODBMS or RDBMS) as is well known in the art. It is understood that "server system" and "server" are often used interchangeably herein. Similarly, the database object described herein can be implemented as single databases, a distributed database, a collection of distributed databases, a database with redundant online or offline backups or other redundancies, etc., and might include a distributed database or storage network and associated processing intelligence.

FIG. 7B illustrates another block diagram of an embodiment of elements of FIG. 7A and various possible interconnections between such elements in accordance with the described embodiments. FIG. 7B also illustrates environment 799. However, in FIG. 7B, the elements of system 716 and various interconnections in an embodiment are illustrated in further detail. More particularly, FIG. 7B shows that user system 712 may include a processor system 712A, memory system 712B, input system 712C, and output system 712D. FIG. 7B shows network 714 and system 716. FIG. 7B also shows that system 716 may include tenant data storage 722, having therein tenant data 723, which includes, for example, tenant storage space 727, tenant data 729, and application metadata 731. System data storage 724 is depicted as having therein system data 725. Further depicted within the expanded detail of application servers 700_{1-N} are User Interface (UI) 730, Application Program Interface (API) 732, application platform 718 includes PL/SOQL 734, save routines 736, application setup mechanism 738, process space 728 includes system process space 702, tenant 1-N process spaces 704, and tenant management process space 710. In other embodiments, environment 799 may not have the same elements as those listed above and/or may have other elements instead of, or in addition to, those listed above.

User system 712, network 714, system 716, tenant data storage 722, and system data storage 724 were discussed above in FIG. 7A. As shown by FIG. 7B, system 716 may

include a network interface 720 (of FIG. 7A) implemented as a set of HTTP application servers 700, an application platform 718, tenant data storage 722, and system data storage 724. Also shown is system process space 702, including individual tenant process spaces 704 and a tenant management process space 710. Each application server 700 may be configured to tenant data storage 722 and the tenant data 723 therein, and system data storage 724 and the system data 725 therein to serve requests of user systems 712. The tenant data 723 might be divided into individual tenant storage areas (e.g., tenant storage space 727), which can be either a physical arrangement and/or a logical arrangement of data. Within each tenant storage space 727, tenant data 729, and application metadata 731 might be similarly allocated for each user. For example, a copy of a user's most recently used (MRU) items might be stored to tenant data 729. Similarly, a copy of MRU items for an entire organization that is a tenant might be stored to tenant storage space 727. A UI 730 provides a user interface and an API 732 provides an application programmer interface into system 716 resident processes to users and/or developers at user systems 712. The tenant data and the system data may be stored in various databases, such as one or more Oracle™ databases.

Application platform 718 includes an application setup mechanism 738 that supports application developers' creation and management of applications, which may be saved as metadata into tenant data storage 722 by save routines 736 for execution by subscribers as one or more tenant process spaces 704 managed by tenant management process space 710 for example. Invocations to such applications may be coded using PL/SOQL 734 that provides a programming language style interface extension to API 732. Invocations to applications may be detected by one or more system processes, which manages retrieving application metadata 731 for the subscriber making the invocation and executing the metadata as an application in a virtual machine.

Each application server 700 may be communicably coupled to database systems, e.g., having access to system data 725 and tenant data 723, via a different network connection. For example, one application server 700₁ might be coupled via the network 714 (e.g., the Internet), another application server 700_{N-1} might be coupled via a direct network link, and another application server 700_N might be coupled by yet a different network connection. Transfer Control Protocol and Internet Protocol (TCP/IP) are typical protocols for communicating between application servers 700 and the database system. However, it will be apparent to one skilled in the art that other transport protocols may be used to optimize the system depending on the network interconnect used.

In certain embodiments, each application server 700 is configured to handle requests for any user associated with any organization that is a tenant. Because it is desirable to be able to add and remove application servers from the server pool at any time for any reason, there is preferably no server affinity for a user and/or organization to a specific application server 700. In one embodiment, therefore, an interface system implementing a load balancing function (e.g., an F5 Big-IP load balancer) is communicably coupled between the application servers 700 and the user systems 712 to distribute requests to the application servers 700. In one embodiment, the load balancer uses a least connections algorithm to route user requests to the application servers 700. Other examples of load balancing algorithms, such as round robin and observed response time, also can be used. For example, in certain embodiments, three consecutive requests from the

same user may hit three different application servers 700, and three requests from different users may hit the same application server 700. In this manner, system 716 is multi-tenant, in which system 716 handles storage of, and access to, different objects, data and applications across disparate users and organizations.

As an example of storage, one tenant might be a company that employs a sales force where each salesperson uses system 716 to manage their sales process. Thus, a user might maintain contact data, leads data, customer follow-up data, performance data, goals and progress data, etc., all applicable to that user's personal sales process (e.g., in tenant data storage 722). In an example of a MTS arrangement, since all of the data and the applications to access, view, modify, report, transmit, calculate, etc., can be maintained and accessed by a user system having nothing more than network access, the user can manage his or her sales efforts and cycles from any of many different user systems. For example, if a salesperson is visiting a customer and the customer has Internet access in their lobby, the salesperson can obtain critical updates as to that customer while waiting for the customer to arrive in the lobby.

While each user's data might be separate from other users' data regardless of the employers of each user, some data might be organization-wide data shared or accessible by a plurality of users or all of the users for a given organization that is a tenant. Thus, there might be some data structures managed by system 716 that are allocated at the tenant level while other data structures might be managed at the user level. Because an MTS might support multiple tenants including possible competitors, the MTS may have security protocols that keep data, applications, and application use separate. Also, because many tenants may opt for access to an MTS rather than maintain their own system, redundancy, up-time, and backup are additional functions that may be implemented in the MTS. In addition to user-specific data and tenant specific data, system 716 might also maintain system level data usable by multiple tenants or other data. Such system level data might include industry reports, news, postings, and the like that are sharable among tenants.

In certain embodiments, user systems 712 (which may be client systems) communicate with application servers 700 to request and update system-level and tenant-level data from system 716 that may require sending one or more queries to tenant data storage 722 and/or system data storage 724. System 716 (e.g., an application server 700 in system 716) automatically generates one or more SQL statements (e.g., one or more SQL queries) that are designed to access the desired information. System data storage 724 may generate query plans to access the requested data from the database.

Each database can generally be viewed as a collection of objects, such as a set of logical tables, containing data fitted into predefined categories. A "table" is one representation of a data object, and may be used herein to simplify the conceptual description of objects and custom objects as described herein. It is understood that "table" and "object" may be used interchangeably herein. Each table generally contains one or more data categories logically arranged as columns or fields in a viewable schema. Each row or record of a table contains an instance of data for each category defined by the fields. For example, a CRM database may include a table that describes a customer with fields for basic contact information such as name, address, phone number, fax number, etc. Another table might describe a purchase order, including fields for information such as customer, product, sale price, date, etc. In some multi-tenant database systems, standard entity tables might be provided for use by

all tenants. For CRM database applications, such standard entities might include tables for Account, Contact, Lead, and Opportunity data, each containing pre-defined fields. It is understood that the word “entity” may also be used interchangeably herein with “object” and “table.”

In some multi-tenant database systems, tenants may be allowed to create and store custom objects, or they may be allowed to customize standard entities or objects, for example by creating custom fields for standard objects, including custom index fields. In certain embodiments, for example, all custom entity data rows are stored in a single multi-tenant physical table, which may contain multiple logical tables per organization. It is transparent to customers that their multiple “tables” are in fact stored in one large table or that their data may be stored in the same table as the data of other customers.

FIG. 8 illustrates a diagrammatic representation of a machine 800 in the exemplary form of a computer system, in accordance with one embodiment, within which a set of instructions, for causing the machine/computer system 800 to perform any one or more of the methodologies discussed herein, may be executed. In alternative embodiments, the machine may be connected (e.g., networked) to other machines in a Local Area Network (LAN), an intranet, an extranet, or the public Internet. The machine may operate in the capacity of a server or a client machine in a client-server network environment, as a peer machine in a peer-to-peer (or distributed) network environment, as a server or series of servers within an on-demand service environment. Certain embodiments of the machine may be in the form of a personal computer (PC), a tablet PC, a set-top box (STB), a Personal Digital Assistant (PDA), a cellular telephone, a web appliance, a server, a network router, switch or bridge, computing system, or any machine capable of executing a set of instructions (sequential or otherwise) that specify actions to be taken by that machine. Further, while only a single machine is illustrated, the term “machine” shall also be taken to include any collection of machines (e.g., computers) that individually or jointly execute a set (or multiple sets) of instructions to perform any one or more of the methodologies discussed herein.

The exemplary computer system 800 includes a processor 802, a main memory 804 (e.g., read-only memory (ROM), flash memory, dynamic random access memory (DRAM) such as synchronous DRAM (SDRAM) or Rambus DRAM (RDRAM), etc., static memory such as flash memory, static random access memory (SRAM), volatile but high-data rate RAM, etc.), and a secondary memory 818 (e.g., a persistent storage device including hard disk drives and a persistent database and/or a multi-tenant database implementation), which communicate with each other via a bus 830. Main memory 804 includes a historical record viewer engine 824 having display logic and functionality to present the historical state of records and associated transactions for such records to a user interface. The corruption UI engine 823 also of main memory 804 presents a user interface through which a user can modify selected records to correct or otherwise rectify corruption in the LSM database. The corruption editor 825 generates compensating transactions and if necessary replacement or corrected extents to the LSM database on behalf of a user based on the user’s selections and modifications at a UI. Main memory 804 and its sub-elements are operable in conjunction with processing logic 826 and processor 802 to perform the methodologies discussed herein. The computer system 800 may additionally or alternatively embody the server side elements as described above.

Processor 802 represents one or more general-purpose processing devices such as a microprocessor, central processing unit, or the like. More particularly, the processor 802 may be a complex instruction set computing (CISC) microprocessor, reduced instruction set computing (RISC) microprocessor, very long instruction word (VLIW) microprocessor, processor implementing other instruction sets, or processors implementing a combination of instruction sets. Processor 802 may also be one or more special-purpose processing devices such as an application specific integrated circuit (ASIC), a field programmable gate array (FPGA), a digital signal processor (DSP), network processor, or the like. Processor 802 is configured to execute the processing logic 826 for performing the operations and functionality which is discussed herein.

The computer system 800 may further include a network interface card 808. The computer system 800 also may include a user interface 810 (such as a video display unit, a liquid crystal display (LCD), or a cathode ray tube (CRT)), an alphanumeric input device 812 (e.g., a keyboard), a cursor control device 814 (e.g., a mouse), and a signal generation device 816 (e.g., an integrated speaker). The computer system 800 may further include peripheral device 836 (e.g., wireless or wired communication devices, memory devices, storage devices, audio processing devices, video processing devices, etc.).

The secondary memory 818 may include a non-transitory machine-readable storage medium or a non-transitory computer readable storage medium or a non-transitory machine-accessible storage medium 831 on which is stored one or more sets of instructions (e.g., software 822) embodying any one or more of the methodologies or functions described herein. The software 822 may also reside, completely or at least partially, within the main memory 804 and/or within the processor 802 during execution thereof by the computer system 800, the main memory 804 and the processor 802 also constituting machine-readable storage media. The software 822 may further be transmitted or received over a network 820 via the network interface card 808.

FIG. 9 shows a diagrammatic representation of a system 900 in which embodiments may operate, be installed, integrated, or configured.

In accordance with one embodiment, there is a system 900 having at least a processor 990 and a memory 995 therein to execute implementing logic and/or instructions 996. According to such an embodiment, the system 900 further includes an LSM database 950 to store records 998; a query interface 935 to process transactions 997 to the LSM database 950 by writing the transactions into immutable extents 951 for subsequent merge into the LSM database 950; a log 999 to record a history of the transactions 997; a corruption locator 930 to identify corruption in at least one of the records 998 of the LSM database 950 represented within the immutable extents 951 prior to merge of the immutable extents 951 to the LSM database 950; and a corruption editor 931 to correct the identified corruption to the at least one of the records 998 of the LSM database 950 based at least in part on the recorded history of the transactions in the log 999 during live operation of the LSM database 950.

According to another embodiment, the system 900 further includes a historical record viewer engine 920 to present a historical records view to a user interface based on the log 999, the historical records view to display one or more records updated during a time range configurable at the user interface.

According to another embodiment, the system 900 further includes a corruption user interface engine 925 to receive

user input selecting one or more of the records updated during the time range; and in which the corruption user interface engine **925** is to further receive user input specifying corruption edits to the one or more records.

According to another embodiment, the system **900** further includes a web-server **925** to implement the historical record view engine and to receive the user input at the system from a client device operating at one of a plurality of customer organizations remote from the system; in which the plurality of customer organizations communicably interface with the system **900** via a public Internet; and in which each customer organization is an entity selected from the group consisting of: a separate and distinct remote organization, an organizational group within the host organization, a business partner of the host organization, or a customer organization that subscribes to cloud computing services provided by the host organization.

According to another embodiment of the system **900**, a user interface operates at a user client device remote from the system and communicatively interfaces with the system via a public Internet; in which the system operates at a host organization as a cloud based service provider to the user client device; and in which the cloud based service provider hosts the customer data within a multi-tenant database system at a host organization on behalf of the user.

Bus **915** interfaces the various components of the system **900** amongst each other, with any other peripheral(s) of the system **900**, and with external components such as external network elements, other machines, client devices, etc., including communicating with such external devices via a network interface over a LAN, WAN, or the public Internet.

According to another embodiment, the system **900** is embodied within one of a tablet computing device or a hand-held smartphone such as those depicted at FIGS. **5A** and **5B**.

While the subject matter disclosed herein has been described by way of example and in terms of the specific embodiments, it is to be understood that the claimed embodiments are not limited to the explicitly enumerated embodiments disclosed. To the contrary, the disclosure is intended to cover various modifications and similar arrangements as are apparent to those skilled in the art. Therefore, the scope of the appended claims are to be accorded the broadest interpretation so as to encompass all such modifications and similar arrangements. It is to be understood that the above description is intended to be illustrative, and not restrictive. Many other embodiments will be apparent to those of skill in the art upon reading and understanding the above description. The scope of the disclosed subject matter is therefore to be determined in reference to the appended claims, along with the full scope of equivalents to which such claims are entitled.

What is claimed is:

1. A method in a host organization having a processor and a memory therein, wherein the method comprises:

storing records in a LSM database at the host organization;

processing transactions to the LSM database by writing the transactions into immutable extents for subsequent merge into the LSM database;

recording a history of the transactions in a log;

identifying corruption in at least one of the records of the LSM database represented within the immutable extents prior to merge of the immutable extents to the LSM database; and

correcting the identified corruption to the at least one of the records of the LSM database based at least in part

on the recorded history of the transactions in the log during live operation of the LSM database.

2. The method of claim **1**, wherein correcting the identified corruption to the at least one of the records of the LSM database comprises at least one of:

restoring a modified record to a prior state;

deleting an erroneously inserted record;

restoring an erroneously deleted record; and

changing one or more field values of an erroneously modified record.

3. The method of claim **1**, wherein correcting the identified corruption to the at least one of the records of the LSM database comprises:

reading a stored record from the LSM database having a corresponding transaction pending merge processing within the immutable extents;

reading compensating transaction data from the log;

generating a compensating transaction based on the compensating transaction data from the log and the stored record from the LSM database; and

sending the compensating transaction to the LSM database.

4. The method of claim **3**, wherein sending the compensating transaction to the LSM database comprises:

issuing the compensating transaction to the LSM database as a new transaction, wherein the new transaction is enqueued within the immutable extent of the LSM database for subsequent merge processing; and

wherein the subsequent merge processing counteracts the corresponding transaction pending merge processing within the immutable extents for the stored record read from the LSM database.

5. The method of claim **3**, wherein sending the compensating transaction to the LSM database comprises:

generating a replacement extent by replaying a plurality of transactions from the log for one of the immutable extents of the LSM database awaiting subsequent merge and updating the identified corruption in the replacement extent using the compensating transaction; and

replacing the one immutable extent for which the plurality of transactions were replayed with the replacement extent.

6. The method of claim **1**, wherein identifying corruption in at least one of the records of the LSM database comprises:

triggering error messages to users for any queries to the LSM database having the identified corruption within the query claim scope; and

maintaining availability and accessibility to the LSM database to database queries.

7. The method of claim **1**, further comprising:

presenting a historical records view to a user interface based on the log, the historical records view displaying one or more records updated during a time range configurable at the user interface.

8. The method of claim **7**, further comprising:

receiving user input selecting one or more of the records updated during the time range; and

receiving user input specifying corruption edits to the one or more records.

9. The method of claim **8**, wherein the user input specifying corruption edits to the one or more records comprises user input specifying one or more of:

a selected record erroneously inserted into the LSM database to be deleted from the LSM database via a compensating delete transaction;

27

a selected record erroneously deleted from the LSM database to be restored to the LSM database via a compensating insert transaction;

a selected record erroneously modified within the LSM database to be reverted to a prior state via a compensating update transaction; and

a selected record erroneously modified within the LSM database to be updated to a new state via a compensating update transaction.

10. The method of claim 7, further comprising:
presenting a preview of other records impacted by a proposed compensating transaction to the user interface prior to proceeding with the proposed compensating transaction.

11. The method of claim 1, further comprising:
receiving a selection from the user interface specifying one of the displayed records updated during the time range;

presenting a record editor to the user interface having the record corresponding to the selection populated therein;

receiving user input to the fields of the record corresponding to the selection from the record editor at the user interface; and

generating a compensating transaction for the record corresponding to the selection using the user input to the fields of the record.

12. The method of claim 1:
wherein the host organization implements the method via computing architecture of the host organization including at least the processor and the memory;

wherein a user interface operates at a user client device remote from the host organization and communicatively interfaces with the host organization via a public Internet; and

wherein the host organization operates as a cloud based service provider to the user client device.

13. The method of claim 11, wherein the host organization provides a multi-tenant database system via the LSM database and the computing architecture of the host organization, the multi-tenant database system having elements of hardware and software that are shared by a plurality of separate and distinct customer organizations, each of the separate and distinct customer organizations being remotely located from the host organization.

14. The method of claim 1, wherein identifying corruption in at least one of the records of the LSM database comprises one of:

identifying corruption responsive to a physical corruption error triggered by the LSM database; and

identifying corruption responsive user input at a corruption user interface, the user input indicating a selected record is to be corrected via a compensating transaction.

15. Non-transitory computer readable storage media having instructions stored thereon that, when executed by a processor in a host organization, the instructions cause the host organization to perform operations comprising:

storing records in a LSM database at the host organization;

processing transactions to the LSM database by writing the transactions into immutable extents for subsequent merge into the LSM database;

recording a history of the transactions in a log;

28

identifying corruption in at least one of the records of the LSM database represented within the immutable extents prior to merge of the immutable extents to the LSM database; and

correcting the identified corruption to the at least one of the records of the LSM database based at least in part on the recorded history of the transactions in the log during live operation of the LSM database.

16. The non-transitory computer readable storage media of claim 15, wherein correcting the identified corruption to the at least one of the records of the LSM database comprises at least one of:

restoring a modified record to a prior state;

deleting an erroneously inserted record;

restoring an erroneously deleted record; and

changing one or more field values of an erroneously modified record.

17. The non-transitory computer readable storage media of claim 15, wherein correcting the identified corruption to the at least one of the records of the LSM database comprises:

reading a stored record from the LSM database having a corresponding transaction pending merge processing within the immutable extents;

reading compensating transaction data from the log;

generating a compensating transaction based on the compensating transaction data from the log and the stored record from the LSM database; and

sending the compensating transaction to the LSM database.

18. The non-transitory computer readable storage media of claim 15, wherein the instructions cause the processor to perform further operations comprising:

presenting a historical records view to a user interface based on the log, the historical records view displaying one or more records updated during a time range configurable at the user interface;

receiving user input selecting one or more of the records updated during the time range;

receiving user input specifying corruption edits to the one or more records, wherein the user input specifying corruption edits to the one or more records comprises user input specifying one or more of:

a selected record erroneously inserted into the LSM database to be deleted from the LSM database via a compensating delete transaction;

a selected record erroneously deleted from the LSM database to be restored to the LSM database via a compensating insert transaction;

a selected record erroneously modified within the LSM database to be reverted to a prior state via a compensating update transaction; and

a selected record erroneously modified within the LSM database to be updated to a new state via a compensating update transaction.

19. A system comprising:

a processor and a memory to execute instructions at the system;

an LSM database to store records;

a query interface to process transactions to the LSM database by writing the transactions into immutable extents for subsequent merge into the LSM database;

a log to record a history of the transactions;

a corruption locator to identify corruption in at least one of the records of the LSM database represented within the immutable extents prior to merge of the immutable extents to the LSM database; and

29

a corruption editor to correct the identified corruption to the at least one of the records of the LSM database based at least in part on the recorded history of the transactions in the log during live operation of the LSM database.

20. The system of claim **19**, further comprising:

a historical record viewer engine to present a historical records view to a user interface based on the log, the historical records view to display one or more records updated during a time range configurable at the user interface.

21. The system of claim **20**, further comprising:

a corruption user interface engine to receive user input selecting one or more of the records updated during the time range; and

the corruption user interface engine to further receive user input specifying corruption edits to the one or more records.

22. The system of claim **21**, further comprising:

a web-server to implement the historical record viewer engine and to receive the user input at the system from a client device operating at one of a plurality of customer organizations remote from the system;

30

wherein the plurality of customer organizations communicably interface with the system via a public Internet; and

wherein each customer organization is an entity selected from the group consisting of: a separate and distinct remote organization, an organizational group within the host organization, a business partner of the host organization, or a customer organization that subscribes to cloud computing services provided by the host organization.

23. The system of claim **19**:

wherein a user interface operates at a user client device remote from the system and communicatively interfaces with the system via a public Internet;

wherein the system operates at a host organization as a cloud based service provider to the user client device; and

wherein the cloud based service provider hosts the customer data within a multi-tenant database system at a host organization on behalf of the user.

* * * * *